

The syslog-ng Store Box 3 LTS Administrator Guide

Publication date May 12, 2015

Abstract

This document is the primary manual of the syslog-ng Store Box 3 LTS.





Copyright © 1996-2015 BalaBit S.a.r.l.

This guide is published under the Creative Commons Attribution-Noncommercial-No Derivative Works (by-nc-nd) 3.0 license. See *Appendix F, Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd) License (p. 218)* for details. The latest version is always available at <https://www.balabit.com/support/documentation>.

This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>). This product includes cryptographic software written by Eric Young (eay@cryptsoft.com)

This documentation and the product it describes are considered protected by copyright according to the applicable laws.

The Zorp™ name and the Zorp™ logo are registered trademarks of BalaBit.

The BalaBit Shell Control Box™ name and the BalaBit Shell Control Box™ logo are registered trademarks of BalaBit.

The syslog-ng™ name and the syslog-ng™ logo are registered trademarks of BalaBit.

The BalaBit™ name and the BalaBit™ logo are registered trademarks of BalaBit.

Linux™ is a registered trademark of Linus Torvalds.

Debian™ is a registered trademark of Software in the Public Interest Inc.

Windows™ 95, 98, ME, 2000, XP, Server 2003, Vista, Server 2008 and 7 are registered trademarks of Microsoft Corporation.

MySQL™ is a registered trademark of MySQL AB in the United States, the European Union and other countries.

Oracle™, JD Edwards™, PeopleSoft™, and Siebel™ are registered trademarks of Oracle Corporation and/or its affiliates.

Sun™, Sun Microsystems™, the Sun logo, Sun Fire 4140™, Sun Fire 2100™, Sun Fire 2200™, Sun Fire 4540™, and Sun StorageTek™ are trademarks or registered trademarks of Sun Microsystems, Inc. or its subsidiaries in the U.S. and other countries.

VMware™, VMware ESX™ and VMware View™ are trademarks or registered trademarks of VMware, Inc. and/or its affiliates.

Citrix®, ICA® and XenApp™ are trademarks or registered trademarks of Citrix Systems, Inc.

All other product names mentioned herein are the trademarks of their respective owners.

Some rights reserved.

DISCLAIMER

BalaBit is not responsible for any third-party Web sites mentioned in this document. BalaBit does not endorse and is not responsible or liable for any content, advertising, products, or other material on or available from such sites or resources. BalaBit will not be responsible or liable for any damage or loss caused or alleged to be caused by or in connection with use of or reliance on any such content, goods, or services that are available on or through any such sites or resources.



Table of Contents

Preface	xi
1. Summary of contents	xi
2. Target audience and prerequisites	xii
3. Products covered in this guide	xii
4. Typographical conventions	xii
5. Contact and support information	xiii
5.1. Sales contact	xiii
5.2. Support contact	xiv
5.3. Training	xiv
6. About this document	xiv
6.1. Summary of changes	xiv
6.2. Feedback	xvi
1. Introduction	1
1.1. What SSB is	1
1.2. What SSB is not	1
1.3. Why is SSB needed	2
1.4. Who uses SSB	2
2. The concepts of SSB	3
2.1. The philosophy of SSB	3
2.2. Collecting logs with SSB	4
2.3. Managing incoming and outgoing messages with flow-control	5
2.3.1. Flow-control and multiple destinations	7
2.4. Receiving logs from a secure channel	7
2.5. Network interfaces	8
2.6. High Availability support in SSB	8
2.7. Firmware in SSB	9
2.7.1. Firmwares and high availability	9
2.8. Versions and releases of SSB	9
2.9. Licenses	10
2.10. The structure of a log message	11
2.10.1. BSD-syslog or legacy-syslog messages	11
2.10.2. IETF-syslog messages	13
3. The Welcome Wizard and the first login	17
3.1. The initial connection to SSB	17
3.1.1. Creating an alias IP address (Microsoft Windows)	18
3.1.2. Creating an alias IP address (Linux)	21
3.1.3. Modifying the IP address of SSB	22
3.2. Configuring SSB with the Welcome Wizard	23
3.3. Configuring storage access in SSB	31
4. Basic settings	33
4.1. Supported web browsers and operating systems	33
4.2. The structure of the web interface	33
4.2.1. Elements of the main workspace	36
4.2.2. Multiple web users and locking	37
4.3. Network settings	37



4.3.1. Configuring the management interface	39
4.3.2. Routing management traffic to the management interface	40
4.4. Date and time configuration	41
4.4.1. Configuring a time (NTP) server	42
4.5. SNMP and e-mail alerts	42
4.5.1. Configuring e-mail alerts	43
4.5.2. Configuring SNMP alerts	44
4.5.3. Querying SSB status information using agents	46
4.6. Configuring system monitoring on SSB	47
4.6.1. Configuring monitoring	48
4.6.2. Health monitoring	49
4.6.3. Preventing disk space fill up	49
4.6.4. Configuring message rate alerting	50
4.6.5. System related traps	52
4.6.6. Alerts related to syslog-ng	53
4.7. Data and configuration archiving and backups	54
4.7.1. Creating configuration and data backups	54
4.7.2. Archiving the collected data	59
4.7.3. Encrypting configuration backups with GPG	62
4.7.4. Parameters of the backup protocols	63
4.7.5. Ownership of the backup files	68
5. User management and access control	69
5.1. Managing SSB users locally	69
5.2. Setting password policies for local users	70
5.3. Managing local usergroups	72
5.4. Managing SSB users from an LDAP database	73
5.5. Authenticating users to a RADIUS server	76
5.6. Managing user rights and usergroups	77
5.6.1. Modifying group privileges	78
5.6.2. Creating new usergroups for the SSB web interface	79
5.6.3. Finding specific usergroups	80
5.6.4. How to use usergroups	80
5.6.5. Built-in usergroups of SSB	81
5.7. Listing and searching configuration changes	83
6. Managing SSB	84
6.1. Controlling SSB — restart, shutdown	84
6.2. Managing a high availability SSB cluster	84
6.2.1. High Availability status explained	86
6.2.2. Adjusting the synchronization speed of DRBD	88
6.2.3. Configuring redundant Heartbeat interfaces	89
6.2.4. Redundant Heartbeat interface status explained	90
6.2.5. Configuring next-hop router monitoring	91
6.3. Upgrading SSB	92
6.3.1. Updating SSB and managing the firmware	93
6.3.2. Upgrading both the core and the boot firmware of a high availability system	95
6.3.3. Reverting to an older firmware version	95
6.3.4. Updating the SSB license	96
6.3.5. Exporting the configuration of SSB	97



6.3.6. Importing the configuration of SSB	98
6.4. Accessing the SSB console	99
6.4.1. Using the console menu of SSB	100
6.5. Enabling SSH access to the SSB host	101
6.6. Changing the root password of SSB	102
6.7. Sealed mode	103
6.7.1. Disabling sealed mode	104
6.8. Out-of-band management of SSB	104
6.9. Managing the certificates used on SSB	105
6.9.1. Generating certificates for SSB	107
6.9.2. Uploading external certificates to SSB	108
6.10. Creating hostlist policies	110
6.10.1. Creating hostlists	110
6.10.2. Importing hostlists from files	111
6.11. Managing SAN access in SSB	112
7. Configuring message sources	114
7.1. Default message sources in SSB	114
7.2. Receiving SNMP messages	114
7.3. Creating message sources in SSB	116
8. Storing messages on SSB	119
8.1. Default logspaces in SSB	119
8.2. Configuring the indexer	120
8.3. Using logstores	120
8.3.1. Viewing encrypted logs with logcat	121
8.4. Creating custom message spaces in SSB	121
8.4.1. Creating a new logstore	121
8.4.2. Creating a new text logspace	124
8.5. Managing log spaces	127
8.6. Accessing log files across the network	128
8.6.1. Sharing log files in standalone mode	128
8.6.2. Sharing log files in domain mode	130
8.6.3. Accessing shared files	133
9. Forwarding messages from SSB	135
9.1. Forwarding log messages to SQL databases	135
9.2. SQL templates in SSB	138
9.2.1. The Legacy template	138
9.2.2. The Full template	138
9.2.3. The Custom template	139
9.3. Forwarding log messages to remote servers	139
9.4. Forwarding log messages to SNMP destinations	142
10. Managing log paths	144
10.1. Default logpaths in SSB	144
10.2. Creating new log paths	144
10.3. Filtering messages	146
10.3.1. Modifying messages using rewrite	147
11. Configuring syslog-ng options	149
11.1. General syslog-ng settings	149
11.2. Timestamping configuration on SSB	151



11.3. Using name resolution on SSB	152
11.4. Setting the certificates used in TLS-encrypted log transport	153
12. Browsing log messages and SSB reports	156
12.1. Using the search interface	156
12.1.1. Customizing columns	157
12.1.2. Adding and removing dynamic columns	158
12.2. Changelogs of SSB	159
12.3. Log messages collected on SSB	160
12.3.1. Metadata collected about log messages	161
12.3.2. Using wildcards and boolean search	162
12.3.3. Using and managing search filters	166
12.3.4. Displaying statistics on search results	168
12.3.5. Browsing encrypted log spaces	171
12.4. Configuration changes of syslog-ng peers	175
12.5. Notifications on archiving and backups	176
12.6. Log message alerts	176
12.7. Statistics collection options	177
12.8. Reports	178
12.8.1. Contents of the default reports	180
12.8.2. Configuring custom reports	180
13. Classifying messages with pattern databases	183
13.1. The structure of the pattern database	183
13.2. How pattern matching works	185
13.3. Searching for rulesets	185
13.4. Creating new rulesets and rules	186
13.5. Exporting databases and rulesets	188
13.6. Importing pattern databases	188
13.7. Using pattern parsers	188
13.8. Using parser results in filters and templates	190
13.9. Using the values of pattern parsers in filters and templates	191
14. Troubleshooting SSB	192
14.1. Network troubleshooting	192
14.2. Gathering data about system problems	193
14.3. Viewing logs on SSB	193
14.4. Collecting logs and system information for error reporting	194
14.5. Status history and statistics	196
14.5.1. Displaying custom syslog-ng statistics	197
14.6. Troubleshooting an SSB cluster	197
14.6.1. Recovering SSB if both nodes broke down	197
14.6.2. Recovering from a split brain situation	198
14.6.3. Replacing a node in an SSB HA cluster	200
14.7. Restoring SSB configuration and data	201
14.8. SAN troubleshooting	201
Appendix A. Package contents inventory	203
Appendix B. syslog-ng Store Box Hardware Installation Guide	204
B.1. Installing the SSB hardware	204
B.2. Installing two SSB units in HA mode	205
B.3. Installing a SAN storage module to SSB	206



B.4. Hardware Troubleshooting	207
B.5. Hardware specifications	207
Appendix C. syslog-ng Store Box Software Installation Guide	208
C.1. Installing the SSB software	208
Appendix D. syslog-ng Store Box VMware Installation Guide	210
D.1. Limitations of SSB under VMware	210
D.2. Installing SSB under VMware ESXi	210
D.3. Modifying the virtual disk size under VMware	211
Appendix E. License contract for BalaBit Product	212
Appendix F. Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd) License	218
Glossary	223
Index	227



List of Examples

2.1. Counting log source hosts	10
4.1. Number of hosts and senders	35
4.2. Creating an early time alert	50
4.3. Using the master alert to indicate unexpected events	51
4.4. Configuring NFS on the remote server	68
8.1. Mounting a shared log space using NFS on Linux	134
12.1. Searching exact match	162
12.2. Complex search queries	163
12.3. Combining keywords in search	163
12.4. Using parentheses in search	163
12.5. Using wildcard ? in search	163
12.6. Using wildcard * in search	164
12.7. Using combined wildcards in search	165
12.8. Searching for special characters	165
13.1. Pattern parser syntax	188
13.2. Using the STRING and ESTRING parsers	189



List of Procedures

2.2. Collecting logs with SSB	4
3.1.1. Creating an alias IP address (Microsoft Windows)	18
3.1.2. Creating an alias IP address (Linux)	21
3.1.3. Modifying the IP address of SSB	22
3.2. Configuring SSB with the Welcome Wizard	23
3.3. Configuring storage access in SSB	31
4.3.1. Configuring the management interface	39
4.3.2. Routing management traffic to the management interface	40
4.4.1. Configuring a time (NTP) server	42
4.5.1. Configuring e-mail alerts	43
4.5.2. Configuring SNMP alerts	44
4.5.3. Querying SSB status information using agents	46
4.6.1. Configuring monitoring	48
4.6.3. Preventing disk space fill up	49
4.6.4. Configuring message rate alerting	50
4.7.1. Creating configuration and data backups	54
4.7.2. Archiving the collected data	59
4.7.3. Encrypting configuration backups with GPG	62
4.7.4.1. Configuring Rsync over SSH	63
4.7.4.2. Configuring SMB	64
4.7.4.3. Configuring NFS	66
5.1. Managing SSB users locally	69
5.2. Setting password policies for local users	70
5.3. Managing local usergroups	72
5.4. Managing SSB users from an LDAP database	73
5.5. Authenticating users to a RADIUS server	76
5.6.1. Modifying group privileges	78
5.6.2. Creating new usergroups for the SSB web interface	79
6.2.3. Configuring redundant Heartbeat interfaces	89
6.2.5. Configuring next-hop router monitoring	91
6.3.1. Updating SSB and managing the firmware	93
6.3.2. Upgrading both the core and the boot firmware of a high availability system	95
6.3.3. Reverting to an older firmware version	95
6.3.4. Updating the SSB license	96
6.3.5. Exporting the configuration of SSB	97
6.3.6. Importing the configuration of SSB	98
6.5. Enabling SSH access to the SSB host	101
6.6. Changing the root password of SSB	102
6.7.1. Disabling sealed mode	104
6.9.1. Generating certificates for SSB	107
6.9.2. Uploading external certificates to SSB	108
6.10.1. Creating hostlists	110
6.10.2. Importing hostlists from files	111
7.2. Receiving SNMP messages	114
7.3. Creating message sources in SSB	116



8.4.1. Creating a new logstore	121
8.4.2. Creating a new text logspace	124
8.6.1. Sharing log files in standalone mode	128
8.6.2. Sharing log files in domain mode	130
9.1. Forwarding log messages to SQL databases	135
9.3. Forwarding log messages to remote servers	139
9.4. Forwarding log messages to SNMP destinations	142
10.2. Creating new log paths	144
10.3.1. Modifying messages using rewrite	147
11.4. Setting the certificates used in TLS-encrypted log transport	153
12.1.1. Customizing columns	157
12.1.2. Adding and removing dynamic columns	158
12.3.3.1. Creating and saving filters for later use	166
12.3.4. Displaying statistics on search results	168
12.3.5.1. Assigning decryption keys to a logstore	171
12.3.5.2. Assigning decryption keys to a user account	172
12.3.5.3. Using temporal decryption keys	174
12.8.2. Configuring custom reports	180
13.4. Creating new rulesets and rules	186
13.8. Using parser results in filters and templates	190
14.1. Network troubleshooting	192
14.3. Viewing logs on SSB	193
14.4. Collecting logs and system information for error reporting	194
14.5.1. Displaying custom syslog-ng statistics	197
14.6.1. Recovering SSB if both nodes broke down	197
14.6.2. Recovering from a split brain situation	198
14.6.3. Replacing a node in an SSB HA cluster	200
14.7. Restoring SSB configuration and data	201
14.8. SAN troubleshooting	201
B.1. Installing the SSB hardware	204
B.2. Installing two SSB units in HA mode	205
B.3. Installing a SAN storage module to SSB	206
C.1. Installing the SSB software	208
D.2. Installing SSB under VMware ESXi	210
D.3. Modifying the virtual disk size under VMware	211



Preface

Welcome to the syslog-ng Store Box 3 LTS Administrator Guide!

This document describes how to configure and manage the syslog-ng Store Box (SSB). Background information for the technology and concepts used by the product is also discussed.

1. Summary of contents

Chapter 1, Introduction (p. 1) describes the main functionality and purpose of the syslog-ng Store Box.

Chapter 2, The concepts of SSB (p. 3) discusses the technical concepts and philosophies behind SSB.

Chapter 3, The Welcome Wizard and the first login (p. 17) describes what to do after assembling SSB — it is a step-by-step guide for the initial configuration.

Chapter 4, Basic settings (p. 33) provides detailed description on configuring and managing SSB as a host.

Chapter 5, User management and access control (p. 69) describes how to manage user accounts and privileges.

Chapter 6, Managing SSB (p. 84) explains the basic management tasks of SSB, including the basic control (for example, shutdown or reboot) of the appliance and upgrading.

Chapter 7, Configuring message sources (p. 114) provides description on using the built-in message sources, creating new message sources and receiving SNMP messages.

Chapter 8, Storing messages on SSB (p. 119) describes how to store log messages in log spaces.

Chapter 9, Forwarding messages from SSB (p. 135) explains how to forward log messages to remote destinations.

Chapter 10, Managing log paths (p. 144) discusses the management of log paths.

Chapter 11, Configuring syslog-ng options (p. 149) describes the configuration options of the syslog-ng server running on syslog-ng Store Box.

Chapter 12, Browsing log messages and SSB reports (p. 156) describes how to browse logs, alerts, and reports online on SSB.

Chapter 13, Classifying messages with pattern databases (p. 183) describes how to parse and classify messages using the pattern database.

Chapter 14, Troubleshooting SSB (p. 192) describes troubleshooting and maintenance procedures of syslog-ng Store Box (SSB).

Appendix A, Package contents inventory (p. 203) lists the contents of the package you receive with the syslog-ng Store Box.

Appendix B, syslog-ng Store Box Hardware Installation Guide (p. 204) describes how to set up the syslog-ng Store Box (SSB) hardware.



Appendix D, syslog-ng Store Box VMware Installation Guide (p. 210) describes how to install syslog-ng Store Box (SSB) as a virtual appliance.

Appendix E, License contract for BalaBit Product (p. 212) includes the text of the End User License Agreement applicable to SSB products.

Appendix F, Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd) License (p. 218) includes the text of the Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd) License applicable to The syslog-ng Store Box 3 LTS Administrator Guide.

The *Glossary (p. 223)* provides definitions of important terms used in this guide.

The *Index* provides cross-references to important terms used in this guide.

2. Target audience and prerequisites

This guide is intended for auditors, consultants, and security experts responsible for securing, auditing, and monitoring server administration processes, especially remote server management. It is also useful for IT decision makers looking for a tool to improve the security and auditability of their servers, or to facilitate compliance to the Sarbanes-Oxley (SOX), the Health Insurance Portability and Accountability Act (HIPAA), Basel II, or the Payment Card Industry (PCI) standard.

The following skills and knowledge are necessary for a successful SSB administrator:

- At least basic system administration knowledge.
- An understanding of networks, TCP/IP protocols, and general network terminology.
- An understanding of system logging and the protocols used in remote system logging.
- Familiarity with the concepts of the syslog-ng and the syslog-ng Agent for Windows applications.
- Working knowledge of the UNIX or Linux operating system is not mandatory but highly useful.

3. Products covered in this guide

This guide describes the use of the syslog-ng Store Box version 3 LTS.



Note

Users of the syslog-ng Store Box are entitled to use the syslog-ng Premium Edition application as a log collector agent for SSB. This guide does not cover the installation and configuration of syslog-ng Premium Edition, see [the syslog-ng documentation](#).

4. Typographical conventions

Before you start using this guide, it is important to understand the terms and typographical conventions used in the documentation. For more information on specialized terms and abbreviations used in the documentation, see the Glossary at the end of this document.

The following kinds of text formatting and icons identify special information in the document.

**Tip**

Tips provide best practices and recommendations.

**Note**

Notes provide additional information on a topic, and emphasize important facts and considerations.

**Warning**

Warnings mark situations where loss of data or misconfiguration of the device is possible if the instructions are not obeyed.

Command

Commands you have to execute.

Emphasis

Reference items, additional readings.

/path/to/file

File names.

Parameters

Parameter and attribute names.

Label

GUI output messages or dialog labels.

Menu

A submenu or menu item in the menu bar.

Button

Buttons in dialog windows.

5. Contact and support information

This product is developed and maintained by BalaBit IT Security Ltd. We are located in Budapest, Hungary. Our address is:

BalaBit IT Security Ltd.
 2 Alíz Street
 H-1117 Budapest, Hungary
 Tel: +36 1 398-6700
 Fax: +36 1 208-0875
 E-mail: <info@balabit.com>
 Web: <https://www.balabit.com/>

5.1. Sales contact

You can directly contact us with sales related topics at the e-mail address <sales@balabit.com>, or leave us your contact information and we call you back.



5.2. Support contact

To access the BalaBit Online Support System (BOSS), sign up for an account at [the MyBalaBit page](#) and [request access to the BalaBit Online Support System \(BOSS\)](#). Online support is available 24 hours a day.

BOSS is available only for registered users with a valid support package.

Support e-mail address: <support@balabit.com>.

Support hotline: +36 1 398 6700 (available from 9 AM to 5 PM CET on weekdays)

5.3. Training

BalaBit IT Security Ltd. holds courses on using its products for new and experienced users. For dates, details, and application forms, visit the <https://www.balabit.com/support/trainings/> webpage.

6. About this document

This guide is a work-in-progress document with new versions appearing periodically.

The latest version of this document can be downloaded from the BalaBit website [here](#).

6.1. Summary of changes

6.1.1. Version 3 LTS -

Changes in product:

- *Procedure 14.8, SAN troubleshooting (p. 201)* has been added to the document.
- The Extended schema for SQL destinations has been removed from the product and from the documentation in SSB version 3.0.1.
- Added a warning about a possible problem when archiving to Windows 2008 R2 hosts using the CIFS protocol to *Section 4.7, Data and configuration archiving and backups (p. 54)*.

6.1.2. Version 2 F1 - 3 LTS

Changes in product:

- Description of the **Don't parse messages** has been added to *Procedure 7.3, Creating message sources in SSB (p. 116)*.
- Figures *Configuring SNMP and e-mail alerting; Default message sources in SSB; Configuring syslog-ng options; Configuring persistent name resolution; Configuring TLS settings for syslog-ng; Creating database destinations; Creating server destinations; Displaying search information; Displaying statistics; Creating a new logstore; Creating a new text logspace; Creating new message sources* have been updated.
- Figures *Creating an early time alert; Using the master alert to indicate unexpected events and Modifying messages using rewrite* have been added to the document.



- Alerts *Message rate was outside the specified limits* and *Too many message rate alerts were generated* have been added to Section 4.6.6, *Alerts related to syslog-ng* (p. 53).
- The list of sources in Section 7.1, *Default message sources in SSB* (p. 114) has been updated with the BSD-syslog (legacy TCP) protocol.
- Procedure 4.6.4, *Configuring message rate alerting* (p. 50) has been added to the document.
- Section 12.3.2, *Using wildcards and boolean search* (p. 162) has been added to the document.
- Procedure 10.3.1, *Modifying messages using rewrite* (p. 147) has been added to the document.
- Procedure 4.6.3, *Preventing disk space fill up* (p. 49) has been added to the document.
- Procedure 14.7, *Restoring SSB configuration and data* (p. 201) has been added to the document.
- Section 7.1, *Default message sources in SSB* (p. 114) has been updated with new default ports.
- Section 4.1, *Supported web browsers and operating systems* (p. 33) has been updated with new supported and tested browsers.
- Section 12.7, *Statistics collection options* (p. 177) has been updated.
- Rate limiting has been removed from Procedure 7.3, *Creating message sources in SSB* (p. 116).

Changes in documentation:

- The document chapters have been restructured. Section 1, *Summary of contents* (p. xi) has been updated with the new structure.
- Screenshots in the HTML version of the document have been resized for better visibility.
- The troubleshooting section has been moved to Chapter 14, *Troubleshooting SSB* (p. 192).
- A warning has been added to Procedure 4.5.1, *Configuring e-mail alerts* (p. 43).

6.1.3. Version 2 LTS - 2 F1

Changes in product:

- Procedure 9.4, *Forwarding log messages to SNMP destinations* (p. 142) has been added to the document.
- References to hardware manuals in Appendix *syslog-ng Store Box Hardware Installation Guide* have been updated and corrected.
- Appendix *syslog-ng Store Box VMware Installation Guide* has been added to the document.
- Procedure *Modifying the IP address of SSB* has been added to the document.
- Section *High Availability status explained* has been updated with 'Converted' status.
- Windows 7 has been added to the supported operating systems list in sections *Supported web browsers and operating systems*, *Supported protocols and client applications* and *Viewing session information and replaying audit trails*.
- A note has been added to Section 8.6.3, *Accessing shared files* (p. 133).

Changes in documentation:

- Section *Summary of changes* has been added to the document.
- Procedures in the HTML version of the document appear on separate HTML pages.



- Labels of cross-references pointing to procedure steps have been corrected.
- Procedures have been restructured to facilitate easier understanding.
- Latin abbreviations have been replaced in document with their English equivalents.
- Editorial changes.

6.2. Feedback

Any feedback is greatly appreciated, especially on what else this document should cover. General comments, errors found in the text, and any suggestions about how to improve the documentation is welcome at documentation@balabit.com.



Chapter 1. Introduction

This chapter introduces the syslog-ng Store Box (SSB) in a non-technical manner, discussing how and why is it useful, and what additional benefits it offers to an existing IT infrastructure.

1.1. What SSB is

SSB is a device that collects, processes, stores, monitors, and manages log messages. It is a central logserver appliance that can receive system (syslog and eventlog) log messages and Simple Network Management Protocol (SNMP) messages from your network devices and computers, store them in a trusted and signed logstore, automatically archive and backup the messages, and also classify the messages using artificial ignorance.

The most notable features of SSB are the following:

- Secure log collection using Transport Layer Security (TLS).
- Trusted, encrypted, signed, timestamped storage.
- Ability to collect log messages from a wide range of platforms, including Linux, Unix, BSD, Sun Solaris, HP-UX, IBM AIX, IBM System i, as well as Microsoft Windows XP, Server 2003, Vista, and Server 2008.
- Forwards messages to log analyzing engines.
- Classifies messages using customizable pattern databases for real-time log monitoring, alerting, and artificial ignorance.
- High Availability (HA) support to ensure continuous log collection in business-critical environments.
- Real-time log monitoring and alerting.
- Retrieves group memberships of the administrators from a Lightweight Directory Access Protocol (LDAP) database.
- Strict, yet easily customizable access control to grant users access only to selected log messages

SSB is configured and managed from any modern web browser that supports HTTPS connections, JavaScript, and cookies. Supported browsers: Mozilla Firefox 10 and Microsoft Internet Explorer 8 and 9. Other tested browsers: Mozilla Firefox 3.6 and Google Chrome 17.

1.2. What SSB is not

SSB is not a log analyzing engine, it is able to classify individual log messages using artificial ignorance, much like the popular logcheck application of the Unix world. SSB comes with a built-in database of log message patterns that are considered “normal”. Messages matching these patterns are produced during the legitimate use of the applications (for example sendmail, Postfix, MySQL, and so on), and are unimportant from the log monitoring perspective, while the remaining messages may contain something “interesting”. The administrators can define log patterns on the SSB interface, label matching messages (for example security event, and so on.) and request alerts if a specific pattern is encountered. For thorough log analysis, SSB can also forward the incoming log messages to external log analyzing engines.



1.3. Why is SSB needed

Log messages contain information about the events happening on the hosts. Monitoring system events is essential for security and system health monitoring reasons. A well-established log management solution offers several benefits to an organization. It ensures that computer security records are stored in sufficient detail, and provides a simple way to monitor and review these logs. Routine log reviews and continuous log analysis help to identify security incidents, policy violations, or other operational problems. Logs also often form the base of auditing and forensic analysis, product troubleshooting and support. There are also several laws, regulations and industrial standards that explicitly require the central collection, periodic review, and long-time archiving of log messages. Examples to such regulations are the Sarbanes-Oxley Act (SOX), the Basel II accord, the Health Insurance Portability and Accountability Act (HIPAA), or the Payment Card Industry Data Security Standard (PCI-DSS).

Built around the popular syslog-ng application used by thousands of organizations worldwide, the syslog-ng Store Box (SSB) brings you a powerful, easy to configure appliance to collect and store your logs. Using the features of the latest syslog-ng Premium Edition to their full power, SSB allows you to collect, process, and store log messages from a wide range of platforms and devices.

All data can be stored in encrypted, digitally signed, and optionally timestamped files, preventing any modification or manipulation, satisfying the highest security standards and policy compliance requirements.

1.4. Who uses SSB

SSB is useful for everyone who has to collect, store, and review log messages. In particular, SSB is invaluable for:

- *Central log collection and archiving:* SSB offers a simple, reliable, and convenient way of collecting log messages centrally. It is essentially a high-capacity log server with high availability support. Being able to collect logs from several different platforms makes it easy to integrate into any environment.
- *Secure log transfer and storage:* Log messages often contain sensitive information and also form the base of audit trails for several applications. Preventing eavesdropping during message transfer and unauthorized access once the messages reach the logserver is essential for security and privacy reasons.
- *Policy compliance:* Many organization must comply to regulations like the Sarbanes-Oxley Act (SOX), the Basel II accord, the Health Insurance Portability and Accountability Act (HIPAA), or the Payment Card Industry Data Security Standard (PCI-DSS). These regulations often have explicit or implicit requirements about log management, such as the central collection of log messages, the use of log analysis to prevent and detect security incidents, or guaranteeing the availability of log messages for an extended period of time — up to several years. SSB helps these organizations to comply with these regulations.
- *Automated log monitoring and log preprocessing:* Monitoring log messages is an essential part of system-health monitoring and security incident detection and prevention. SSB offers a powerful platform that can classify tens of thousands of messages real-time to detect messages that deviate from regular messages, and promptly raise alerts. Although this classification does not offer as complete inspection as a log analyzing application, SSB can process much more messages than a regular log analyzing engine, and also filter out unimportant messages to decrease the load on the log analyzing application.

Chapter 2. The concepts of SSB

This chapter discusses the technical concepts of SSB.

2.1. The philosophy of SSB

The syslog-ng Store Box (SSB) is a log server appliance that collects, stores and monitors the log messages sent by network devices, applications and computers. SSB can receive traditional syslog messages, syslog messages that comply with the new Internet Engineering Task Force (IETF) standard, eventlog messages from Microsoft Windows hosts, as well as SNMP messages.

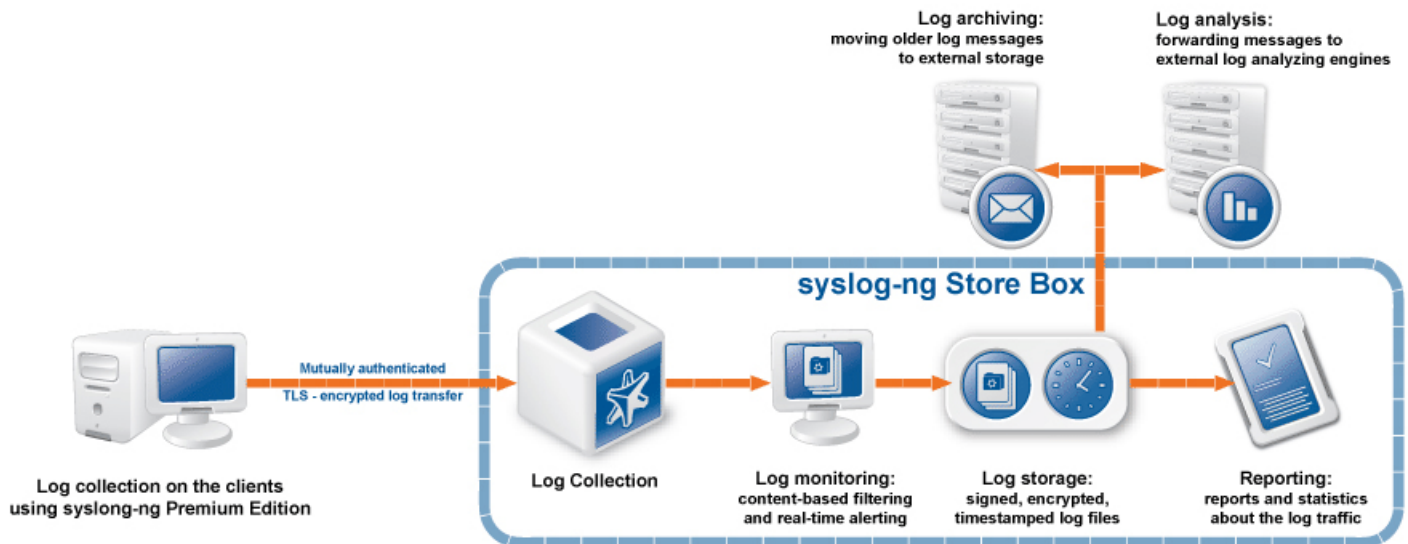


Figure 2.1. The philosophy of the syslog-ng Store Box

Clients can send messages to SSB using their own logging application if it supports the *BSD-syslog* (RFC 3164) or the *IETF-syslog* (RFC 5424-5428) protocol, or they can use the syslog-ng Premium Edition application to act as the log-forwarding agent of SSB.

The main purpose of SSB is to collect the logs from the clients and store them on its hard disk. The messages are stored in so-called logspaces. There are two types of logspaces: the first stores messages in traditional plain-text files, while the second one uses a binary format that can be compressed, encrypted, digitally signed, and also timestamped.

The syslog-ng application reads incoming messages and forwards them to the selected *destinations*. The syslog-ng application can receive messages from files, remote hosts, and other *sources*.

Log messages enter syslog-ng in one of the defined sources, and are sent to one or more *destinations*. In case of the clients, one of the destinations is the syslog-ng Store Box; the destinations on the SSB can be logspaces or remote servers, such as database servers or log analyzing engines.

Sources and destinations are independent objects; *log paths* define what syslog-ng does with a message, connecting the sources to the destinations. A log path consists of one or more sources and one or more destinations; messages arriving to a source are sent to every destination listed in the log path. A log path defined in syslog-ng is called a *log statement*.

Optionally, log paths can include *filters*. Filters are rules that select only certain messages, for example, selecting only messages sent by a specific application. If a log path includes filters, syslog-ng sends only the messages satisfying the filter rules to the destinations set in the log path.

SSB is configured by an administrator or auditor using a web browser.

2.2. Procedure – Collecting logs with SSB

Purpose:

The following procedure illustrates the route of a log message from its source on the syslog-ng client to the syslog-ng Store Box.

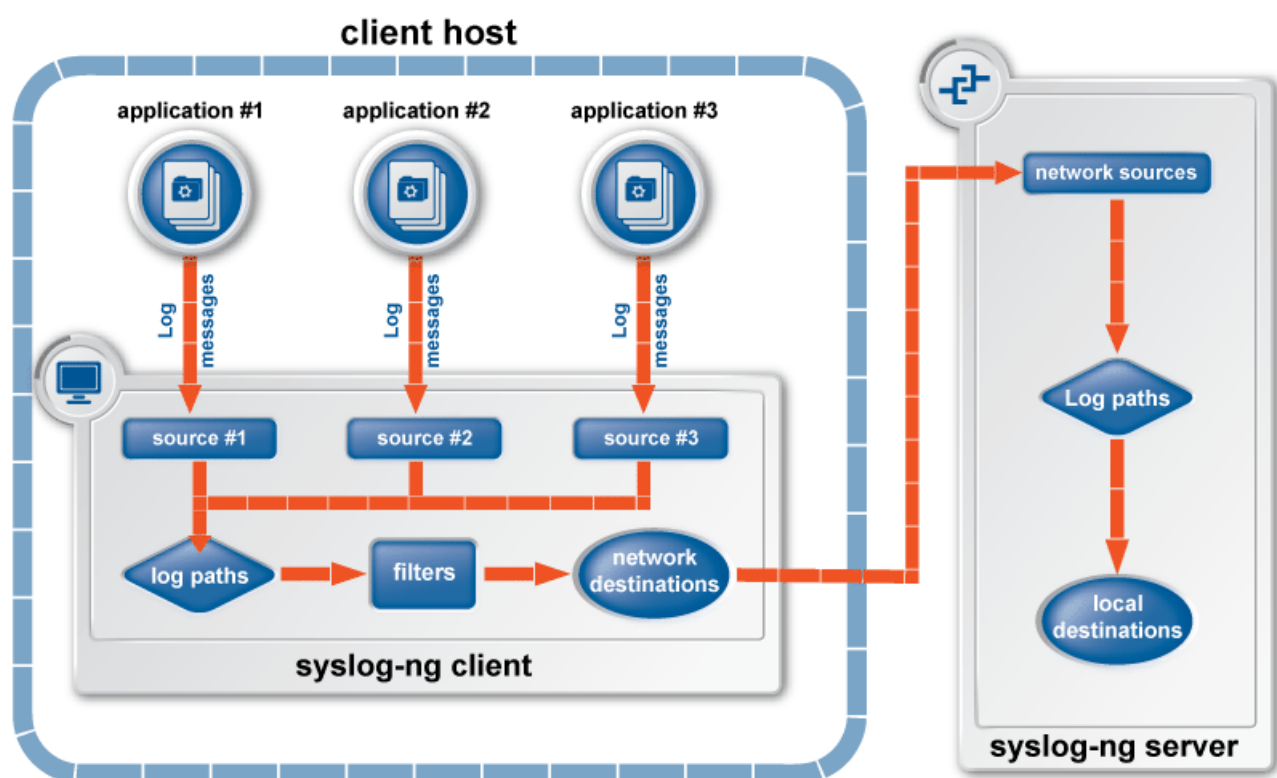


Figure 2.2. The route of a log message

Steps:

- Step 1. A device or application sends a log message to a source on the syslog-ng client. For example, an Apache web server running on Linux enters a message into the `/var/log/apache` file.
- Step 2. The syslog-ng client running on the web server reads the message from its `/var/log/apache` source.
- Step 3. The syslog-ng client processes the first log statement that includes the `/var/log/apache` source.



- Step 4. The syslog-ng client performs optional operations (for example message filtering) on the message; for example, it compares the message to the filters of the log statement (if any). If the message complies with all filter rules, syslog-ng sends the message to the destinations set in the log statement, for example, to the remote syslog-ng server.
After that, the syslog-ng client processes the next log statement that includes the `/var/log/apache` source, repeating Steps 3-4.
- Step 5. The message sent by the syslog-ng client arrives to a source set on the syslog-ng Store Box.
- Step 6. The syslog-ng Store Box reads the message from its source and processes the first log path that includes that source.
- Step 7. The syslog-ng server performs optional operations (for example message filtering, or pattern matching to compare the message to a list of known messages). If the message complies with all filter rules, SSB sends the message to the destinations set in the log path. The destinations are local, optionally encrypted files on SSB, or remote servers such as a database server.
- Step 8. SSB processes the next log statement, repeating Steps 6-8.

**Note**

The syslog-ng application can stop reading messages from its sources if the destinations cannot process the sent messages. This feature is called flow-control and is detailed in *Section 2.3, Managing incoming and outgoing messages with flow-control (p. 5)*.

2.3. Managing incoming and outgoing messages with flow-control

This section describes the internal message-processing model of syslog-ng, as well as the flow-control feature that can prevent message loss. To use flow-control, the **Flow** option must be enabled for the particular log path.

The syslog-ng application monitors (polls) the sources defined in its configuration file, periodically checking each source for messages. When a log message is found in one of the sources, syslog-ng polls every source and reads the available messages. These messages are processed and put into the output buffer of syslog-ng (also called fifo). From the output buffer, the operating system sends the messages to the appropriate destinations.

In large-traffic environments many messages can arrive during a single poll loop, therefore syslog-ng reads only a fixed number of messages from each source. The **Messages fetched in a single poll** option specifies the number of messages read during a poll loop from a single source.




Figure 2.3. Managing log messages in syslog-ng



Note

The **Messages fetched in a single poll** option of SSB can be set as a global option at **Log > Options**.

Every destination has its own output buffer. The output buffer is needed because the destination might not be able to accept all messages immediately. On SSB, the **Output memory buffer** parameter sets the size of the output buffer. The output buffer must be larger than the **Messages fetched in a single poll** of the sources, to ensure that every message read during the poll loop fits into the output buffer. If the log path sends messages to a destination from multiple sources, the output buffer must be large enough to store the incoming messages of every source.

TCP and TLS sources can receive the logs from several incoming connections (for example many different clients or applications). For such sources, syslog-ng reads messages from every connection, thus the **Messages fetched in a single poll** parameter applies individually to every connection of the source.




Figure 2.4. Managing log messages of TCP sources in syslog-ng

The flow-control of syslog-ng introduces a control window to the source that tracks how many messages can syslog-ng accept from the source. Every message that syslog-ng reads from the source decreases the number of free slots by one; every message that syslog-ng successfully sends from the output buffer increases the number of free slots by one. If the window is full (that is, there are no free slots), syslog-ng stops reading messages from the source. The initial size of the control window is by default **100**: the **Output memory buffer** must be larger than this value in order for flow-control to have any effect. If a source accepts messages from multiple connections, all messages use the same control window.

When flow-control is used, every source has its own control window. As a worst-case situation, the output buffer of the destination must be set to accommodate all messages of every control window, that is, the **Output memory buffer** of the destination must be greater than **<Number of sources> * <Initial window size>**. This applies to every source that sends logs to the particular destination, thus if two sources having several connections



and heavy traffic send logs to the same destination, the control windows of every source must fit into the output buffer of the destination. Otherwise, syslog-ng does not activate the flow-control, and messages may be lost.

2.3.1. Flow-control and multiple destinations

Using flow-control on a source has an important side-effect if the messages of the source are sent to multiple destinations. If flow-control is in use and one of the destinations cannot accept the messages, the other destinations do not receive any messages either, because syslog-ng stops reading the source. For example, if messages from a source are sent to a remote server and also stored locally in a file, and the network connection to the server becomes unavailable, neither the remote server nor the local file will receive any messages. This side-effect of the flow-control can be avoided by using the disk-based buffering feature of syslog-ng.



Note

Creating separate log paths for the destinations that use the same flow-controlled source does not help avoiding the problem.

2.4. Receiving logs from a secure channel

The syslog-ng Store Box receive log messages securely over the network using the Transport Layer Security (TLS) protocol (TLS is an encryption protocol over the TCP/IP network protocol).

TLS uses certificates to authenticate and encrypt the communication, as illustrated on the following figure:



Figure 2.5. Certificate-based authentication

The client sending the logs authenticates SSB by requesting its certificate and public key. Optionally, SSB can also request a certificate from the client, thus mutual authentication is also possible.

In order to use TLS encryption in syslog-ng, the following elements are required:

- A certificate on SSB that identifies SSB. This is available by default.
- The certificate of the Certificate Authority that issued the certificate of SSB must be available on the syslog-ng client.



When using mutual authentication to verify the identity of the clients, the following elements are required:

- A certificate must be available on the syslog-ng client. This certificate identifies the syslog-ng client.
- The certificate of the Certificate Authority that issued the certificate of the syslog-ng client must be available on SSB.

Mutual authentication ensures that SSB accepts log messages only from authorized clients.

For details on configuring TLS communication in syslog-ng, see *Chapter 7, Configuring message sources (p. 114)*.

2.5. Network interfaces

The SSB hardware has five network interfaces: the external, the management, the internal currently not used in SSB, the HA, and the IPMI interface. For details on hardware installation, see *Appendix B, syslog-ng Store Box Hardware Installation Guide (p. 204)*.

The *external* interface is used for communication between SSB and the clients: clients send the syslog messages to the external interface of SSB. Also, the initial configuration of SSB is always performed using the external interface (For details on the initial configuration, see *Procedure 3.2, Configuring SSB with the Welcome Wizard (p. 23)*). The external interface is used for management purposes if the management interface is not configured. The external interface uses the Ethernet connector labeled as *EXT*.

The *management* interface is used exclusively for communication between SSB and the auditors or the administrators of SSB. Incoming connections are accepted only to access the SSB web interface, other connections targeting this interface are rejected. The management interface uses the Ethernet connector labeled as *MGMT*.

The routing rules determine which interface is used for transferring remote backups and syslog messages of SSB.



Tip

It is recommended to direct backups, syslog and SNMP messages, and e-mail alerts to the management interface. For details, see *Procedure 4.3.2, Routing management traffic to the management interface (p. 40)*.

If the management interface is not configured, the external interface takes the role of the management interface.

The *HA* interface is an interface reserved for communication between the nodes of SSB clusters. The HA interface uses the Ethernet connector labeled as *HA*. For details on high availability, see *Section 2.6, High Availability support in SSB (p. 8)*.

The Intelligent Platform Management Interface (*IPMI*) interface allows system administrators to monitor system health and to manage SSB events remotely. IPMI operates independently of the operating system of SSB.

2.6. High Availability support in SSB

In high availability (HA) mode two SSB units (called master and slave nodes) having identical configuration are operating simultaneously. The master shares all data with the slave node, and if the master node stops



functioning, the other one becomes immediately active, so the servers are continuously accessible. The slave node takes over the MAC addresses of the interfaces of the master node.



Warning

Hazard of data loss! If the two nodes cannot communicate, the slave assumes that the master node broke down and becomes active, which can lead to both nodes being active at the same time (a split brain situation). This might result in data loss. For details on how to recover from a split brain situation, see *Procedure 14.6.2, Recovering from a split brain situation* (p. 198). For details on how to avoid split brain situations, see *Procedure 6.2.3, Configuring redundant Heartbeat interfaces* (p. 89).

The slave node automatically synchronizes its hard disk with the master via the HA network interface. The disks must be synchronized for the HA support to operate correctly.



Warning

When using the management interface and high availability together, do not forget to connect the management interface of both SSB nodes to the network. Otherwise you will not be able to remotely access SSB if a takeover occurs.

2.7. Firmware in SSB

The SSB firmware is separated into two parts: an *external* and an *internal* firmware.

- The *external* firmware (also called boot firmware) boots up SSB, provides the high availability support, and starts the internal firmware. The external firmware changes very rarely.
- The *internal* firmware (also called core firmware) handles everything else: provides the web interface, receives and processes log messages and so on. The internal firmware is updated regularly as new features are added to SSB.

Both firmwares can be updated from the SSB web interface. For details, see *Section 6.3, Upgrading SSB* (p. 92).

2.7.1. Firmwares and high availability

When powering on the SSB nodes in high availability mode, both nodes boot and start the boot firmware. The boot firmwares then determine which unit is the master: the core firmware is started only on the master node.

Upgrading the SSB firmware via the web interface automatically upgrades the firmware on both nodes.

2.8. Versions and releases of SSB

As of June 2011, the following release policy applies to syslog-ng Store Box:

- *Long Term Supported or LTS releases* (for example, SSB 3 LTS) are supported for 3 years after their original publication date and for 1 year after the next LTS release is published (whichever date is later). The second digit of the revisions of such releases is 0 (for example, SSB 3.0.1). Maintenance releases to LTS releases contain only bugfixes and security updates.
- *Feature releases* (for example, SSB 3 F1) are supported for 6 months after their original publication date and for 2 months after succeeding Feature or LTS Release is published (whichever date is later).



Feature releases contain enhancements and new features, presumably 1-3 new feature per release. Only the last feature release is supported (for example when a new feature release comes out, the last one becomes unsupported within two months).

For a full description on stable and feature releases, see [Stable and feature releases](#).



Warning

Downgrading from a feature release to an earlier (and thus unsupported) feature release, or to the stable release is not supported: this means that once you upgrade a system from a stable release (for example 1.0) to a feature release (for example 1.1), you will have to keep upgrading to the new feature releases until the next stable version release (for example 2.0) is published, or risk using an unsupported product.

2.9. Licenses

SSB's license determines the number of individual hosts (also called log source hosts) that can send log messages to SSB.

A log source host is a host or network device (including syslog-ng clients and relays) that sends logs to the syslog-ng server. Log source hosts can be servers, routers, desktop computers, or other devices capable of sending syslog messages or running syslog-ng. Log source hosts are identified by their IP addresses, so virtual machines and vhosts are separately counted.



Warning

The `chain_hostnames()` option of syslog-ng can interfere with the way syslog-ng counts the log source hosts, causing syslog-ng to act as if there were more hosts logging to the central server. As `chain_hostnames()` is a deprecated option, disable it on your log sources to avoid any problems related to license counting.

The SSB license also allows you to download the syslog-ng Premium Edition application (including the syslog-ng Agent for Windows) and install it on any supported platform to use it as a log collector agent for SSB.

- **syslog-ng Store Box SSB1000 and SSB1000d:** License permits to collect logs from 50 IP addresses. It is upgradable to unlimited.
- **syslog-ng Store Box SSB5000 and SSB10000:** License permits to collect logs from an unlimited number of IP addresses.

Contact BalaBit or your local distributor for details. For details on installing a new license, see *Procedure 6.3.4, Updating the SSB license (p. 96)*.



Example 2.1. Counting log source hosts

In this example you have two facilities (for example data centers or server farms), and you have 80 AIX servers and 20 Microsoft Windows host at Facility 1, and 5 HP-UX servers and 40 Debian servers at Facility 2. That is 145 hosts altogether.

- If you want to collect the log messages of these host to a single SSB, then you need a license that allows you to accept logs from at least 145 hosts. (In practice this means you have to buy a license for 150 hosts.)
- If you want each facility to have its own SSB, and do not want to have a central server that collects the log messages of both facilities, you need two separate SSBs: one with license for 100 hosts at Facility 1, and one with a license for at least 45 hosts at Facility 2 (actually you have to buy license for 50 hosts).
- If you want each facility to have its own local logserver that stores the logs locally, and also want to have a central logserver that collects every log message independently from the two local logserver, you need three SSBs with the following licenses: a license for 100 hosts at Facility 1, and a license for at least 45



hosts at Facility 2, and a license for the central logserver. The size of the license on the central logserver should be 100 (the hosts at Facility 1) + 45 (the hosts at Facility 2) + 2 (the two local logservers at each facility) = 147 — practically that's another 150-host license.



Note

If, for example, the 40 Debian servers at Facility 2 are each running 3 virtual hosts, then the total number of hosts at Facility 2 is 125, and the license sizes should be calculated accordingly.

2.10. The structure of a log message

The following sections describe the structure of log messages. Currently there are two standard syslog message formats:

- The old standard described in RFC 3164 (also called the BSD-syslog or the legacy-syslog protocol): see *Section 2.10.1, BSD-syslog or legacy-syslog messages (p. 11)*
- The new standard described in RFC 5424 (also called the IETF-syslog protocol): see *Section 2.10.2, IETF-syslog messages (p. 13)*

2.10.1. BSD-syslog or legacy-syslog messages

This section describes the format of a syslog message, according to the legacy-syslog or BSD-syslog protocol (see *RFC 3164*). A syslog message consists of the following parts:

- PRI
- HEADER
- MSG

The total message must be shorter than 1024 bytes.

The following is a sample syslog message: `<133>Feb 25 14:09:07 webserver syslogd: restart`. The message corresponds to the following format: `<priority>timestamp hostname application: message`. The different parts of the message are explained in the following sections.



Note

The syslog-ng application supports longer messages as well. For details, see the **Message size** option. However, it is not recommended to enable messages larger than the packet size when using UDP destinations.

2.10.1.1. The PRI message part

The PRI part of the syslog message (known as Priority value) represents the Facility and Severity of the message. Facility represents the part of the system sending the message, while severity marks its importance. The Priority value is calculated by first multiplying the Facility number by 8 and then adding the numerical value of the Severity. The possible facility and severity values are presented below.

**Note**

Facility codes may slightly vary between different platforms. The syslog-ng application accepts facility codes as numerical values as well.

The following table lists the facility values.

Numerical Code	Facility
0	kernel messages
1	user-level messages
2	mail system
3	system daemons
4	security/authorization messages
5	messages generated internally by syslogd
6	line printer subsystem
7	network news subsystem
8	UUCP subsystem
9	clock daemon
10	security/authorization messages
11	FTP daemon
12	NTP subsystem
13	log audit
14	log alert
15	clock daemon
16-23	locally used facilities (local0-local7)

Table 2.1. syslog Message Facilities

The following table lists the severity values.

Numerical Code	Severity
0	Emergency: system is unusable
1	Alert: action must be taken immediately
2	Critical: critical conditions
3	Error: error conditions
4	Warning: warning conditions
5	Notice: normal but significant condition
6	Informational: informational messages



Numerical Code	Severity
7	Debug: debug-level messages

Table 2.2. syslog Message Severities

2.10.1.2. The HEADER message part

The HEADER part contains a timestamp and the hostname (without the domain name) or the IP address of the device. The timestamp field is the local time in the *Mmm dd hh:mm:ss* format, where:

- *Mmm* is the English abbreviation of the month: Jan, Feb, Mar, Apr, May, Jun, Jul, Aug, Sep, Oct, Nov, Dec.
- *dd* is the day of the month on two digits. If the day of the month is less than 10, the first digit is replaced with a space. (For example *Aug 7*.)
- *hh:mm:ss* is the local time. The hour (hh) is represented in a 24-hour format. Valid entries are between 00 and 23, inclusive. The minute (mm) and second (ss) entries are between 00 and 59 inclusive.

2.10.1.3. The MSG message part

The MSG part contains the name of the program or process that generated the message, and the text of the message itself. The MSG part is usually in the following format: *program[pid]: message text*.

2.10.2. IETF-syslog messages

This section describes the format of a syslog message, according to the IETF-syslog protocol (see [RFC 5424-5428](#)). A syslog message consists of the following parts:

- HEADER (includes the PRI as well)
- STRUCTURED-DATA
- MSG

The following is a sample syslog message:

```
<34>1 2003-10-11T22:14:15.003Z mymachine.example.com su - ID47 - BOM'su root' failed
for lonvick on /dev/pts/8
```

The message corresponds to the following format:

```
<priority>VERSION ISOTIMESTAMP HOSTNAME APPLICATION PID MESSAGEID STRUCTURED-DATA
MSG
```

In this example, the Facility has the value of 4, severity is 2, so PRI is 34. The VERSION is 1. The message was created on 11 October 2003 at 10:14:15pm UTC, 3 milliseconds into the next second. The message originated from a host that identifies itself as "mymachine.example.com". The APP-NAME is "su" and the PROCID is unknown. The MSGID is "ID47". The MSG is "'su root' failed for lonvick...", encoded in UTF-8. The encoding is defined by the BOM. There is no STRUCTURED-DATA present in the message, this is indicated by "-" in the STRUCTURED-DATA field. The MSG is "'su root' failed for lonvick...".

Source: <http://tools.ietf.org/html/rfc5424>

The HEADER part of the message must be in plain ASCII format, the parameter values of the STRUCTURED-DATA part must be in UTF-8, while the MSG part should be in UTF-8. The different parts of the message are explained in the following sections.

2.10.2.1. The PRI message part

The PRI part of the syslog message (known as Priority value) represents the Facility and Severity of the message. Facility represents the part of the system sending the message, while severity marks its importance. The Priority value is calculated by first multiplying the Facility number by 8 and then adding the numerical value of the Severity. The possible facility and severity values are presented below.



Note

Facility codes may slightly vary between different platforms. The syslog-ng application accepts facility codes as numerical values as well.

Numerical Code	Facility
0	kernel messages
1	user-level messages
2	mail system
3	system daemons
4	security/authorization messages
5	messages generated internally by syslogd
6	line printer subsystem
7	network news subsystem
8	UUCP subsystem
9	clock daemon
10	security/authorization messages
11	FTP daemon
12	NTP subsystem
13	log audit
14	log alert
15	clock daemon
16-23	locally used facilities (local0-local7)

Table 2.3. syslog Message Facilities

The following table lists the severity values.

Numerical Code	Severity
0	Emergency: system is unusable

Numerical Code	Severity
1	Alert: action must be taken immediately
2	Critical: critical conditions
3	Error: error conditions
4	Warning: warning conditions
5	Notice: normal but significant condition
6	Informational: informational messages
7	Debug: debug-level messages

Table 2.4. syslog Message Severities

2.10.2.2. The HEADER message part

The HEADER part contains the following elements:

- **VERSION**: Version number of the syslog protocol standard. Currently this can only be 1.
- **ISOTIMESTAMP**: The time when the message was generated in the ISO 8601 compatible standard timestamp format (yyyy-mm-ddThh:mm:ss+-ZONE), for example: 2006-06-13T15:58:00.123+01:00.
- **HOSTNAME**: The machine that originally sent the message.
- **APPLICATION**: The device or application that generated the message
- **PID**: The process name or process ID of the syslog application that sent the message. It is not necessarily the process ID of the application that generated the message.
- **MESSAGEID**: The ID number of the message.



Note

The syslog-ng application supports other timestamp formats as well, like ISO, or the PIX extended format. The timestamp used in the IETF-syslog protocol is derived from RFC3339, which is based on ISO8601. For details, see the `ts_format()` option in *The syslog-ng Premium Edition Administrator Guide*.

2.10.2.3. The STRUCTURED-DATA message part

The STRUCTURED-DATA message part may contain meta- information about the syslog message, or application-specific information such as traffic counters or IP addresses. STRUCTURED-DATA consists of data blocks enclosed in brackets (`[]`). Every block include the ID of the block, and one or more *name=value* pairs. The syslog-ng application automatically parses the STRUCTURED-DATA part of syslog messages, which can be referenced in macros (see *The syslog-ng Premium Edition Administrator Guide* for details). An example STRUCTURED-DATA block looks like:

```
[exampleSDID@0 iut="3" eventSource="Application" eventID="1011"] [examplePriority@0 class="high"]
```



2.10.2.4. The MSG message part

The MSG part contains the text of the message itself. The encoding of the text must be UTF-8 if the BOM character is present in the message. If the message does not contain the BOM character, the encoding is treated as unknown. Usually messages arriving from legacy sources do not include the BOM character.



Chapter 3. The Welcome Wizard and the first login

This chapter describes the initial steps of configuring SSB. Before completing the steps below, unpack, assemble, and power on the hardware. Connect at least the external network interface to the local network, or directly to the computer from which SSB will be configured.

**Note**

For details on unpacking and assembling the hardware, see *Appendix B, syslog-ng Store Box Hardware Installation Guide* (p. 204). For details on how to create a high availability SSB cluster, see *Procedure B.2, Installing two SSB units in HA mode* (p. 205).

3.1. The initial connection to SSB

SSB can be connected from a client machine using any modern web browser.

**Note**

For details on supported browsers, see *Section 4.1, Supported web browsers and operating systems* (p. 33)

SSB can be accessed from the local network. Starting with version 2.1, SSB attempts to receive an IP address automatically via DHCP. If it fails to obtain an automatic IP address, it starts listening for HTTPS connections on the `192.168.1.1` IP address. Note that certain switch configurations and security settings can interfere with SSB receiving an IP address via DHCP. SSB accepts connections via its *external* interface (*EXT*, for details on the network interfaces, see *Section 2.5, Network interfaces* (p. 8)).

**Tip**

The SSB console displays the IP address the external interface is listening on.

If SSB is listening on the `192.168.1.1` address, note that the `192.168.1.0/24` subnet must be accessible from the client. If the client machine is in a different subnet (for example its IP address is `192.168.10.X`), but in the same network segment, the easiest way is to assign an alias IP address to the client machine. Creating an alias IP on the client machine virtually puts both the client and SSB into the same subnet, so that they can communicate. To create an alias IP complete the following steps.

- For details on creating an alias IP on Microsoft Windows, see *Procedure 3.1.1, Creating an alias IP address (Microsoft Windows)* (p. 18).
- For details on creating an alias IP on Linux, see *Procedure 3.1.2, Creating an alias IP address (Linux)* (p. 21).



- If configuring an alias interface is not an option for some reason, you can modify the IP address of SSB. For details, see *Procedure 3.1.3, Modifying the IP address of SSB* (p. 22).

**Warning**

The Welcome Wizard can be accessed only using the external network interface of SSB, as the management interface is not configured yet.

3.1.1. Procedure – Creating an alias IP address (Microsoft Windows)

Purpose:

This procedure describes how to assign an alias IP address to a network interface on Microsoft Windows platforms.

Steps:

Step 1. Navigate to **Start menu > Settings > Network Connections**.

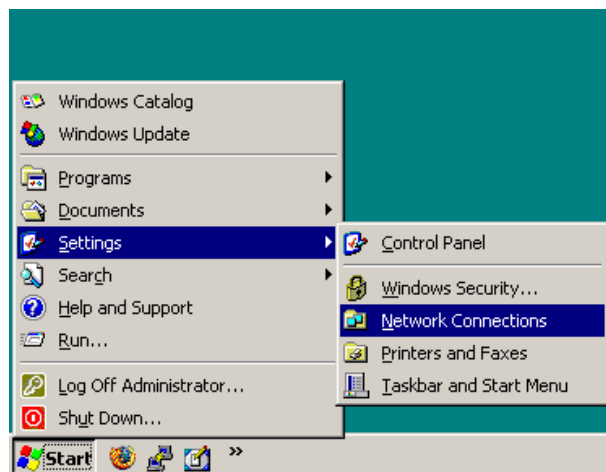


Figure 3.1.

Step 2. Double click on the **Local Area Connection** and then click **Properties**.

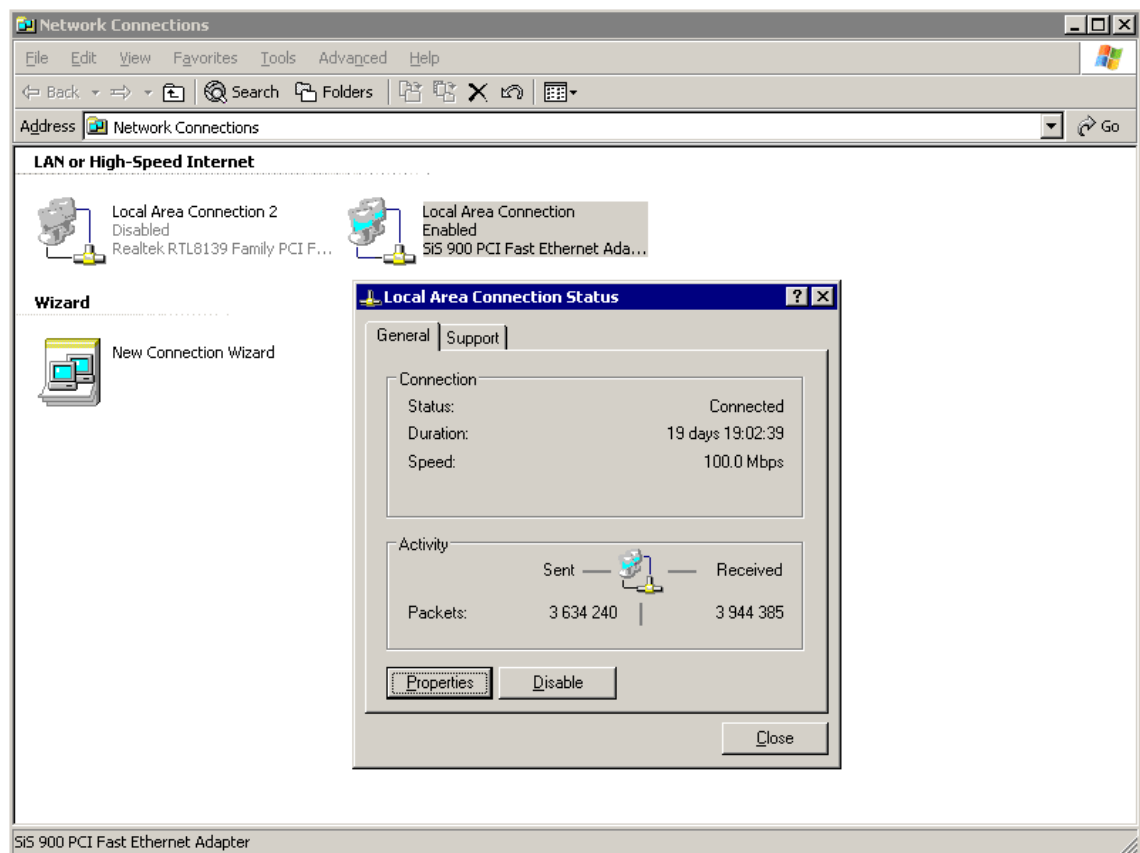


Figure 3.2.

Step 3. Select the **Internet Protocol (TCP/IP)** component in the list and click **Properties**.

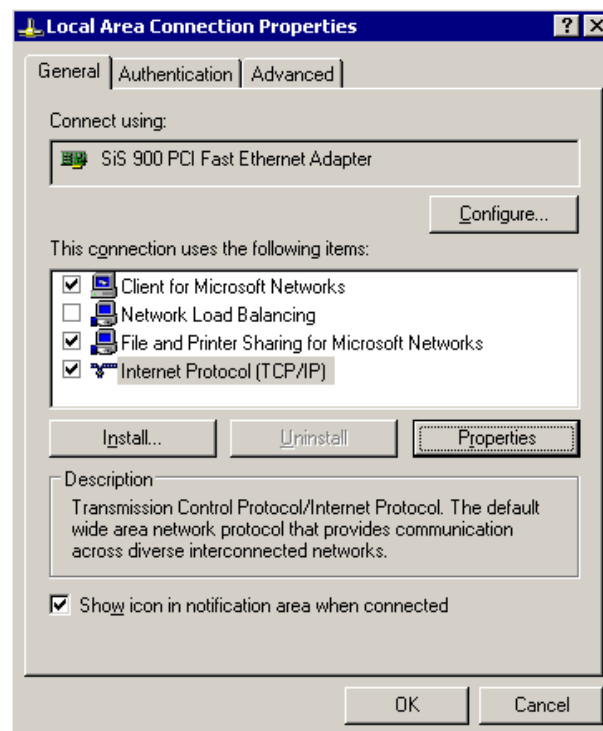


Figure 3.3.

Step 4. To display the Advanced TCP/IP Settings window, click **Advanced**.

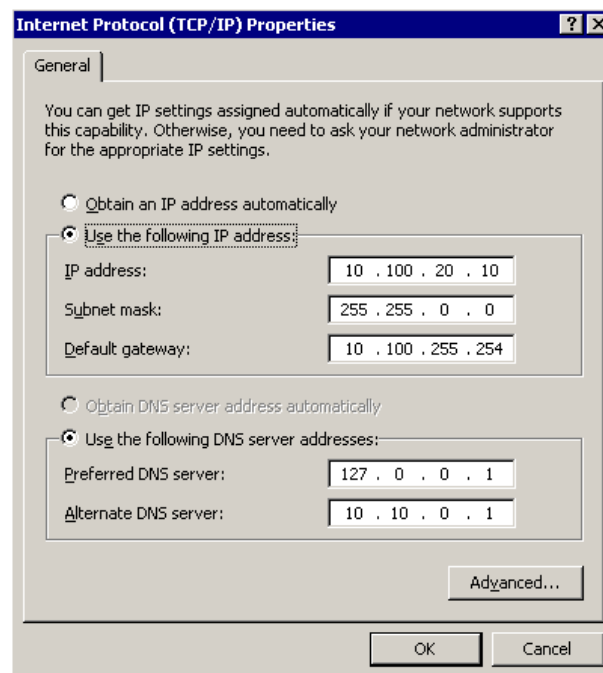


Figure 3.4.

Step 5. Select the **IP Settings** tab and in the **IP Addresses** section, click **Add**.

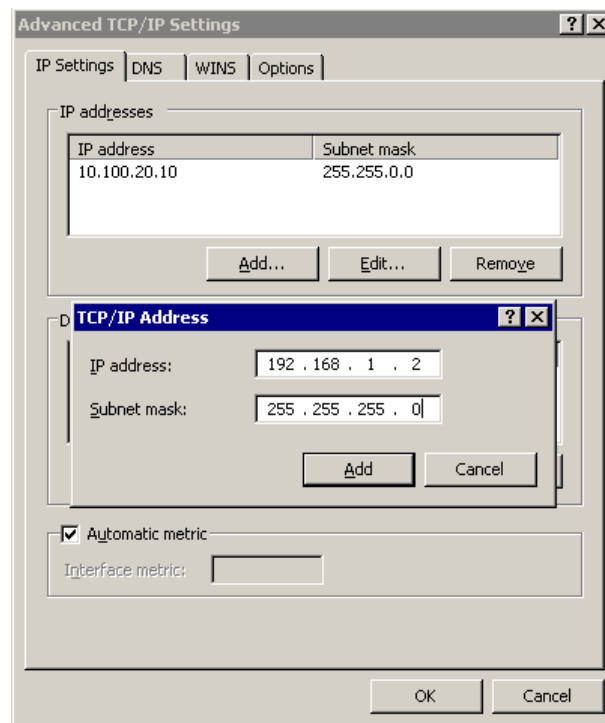


Figure 3.5.

Step 6. Into the **IP Address** field, enter `192.168.1.2`. Into the **Netmask** field, enter `255.255.255.0`.

**Warning**

If your internal network uses the `192.168.1.0/24` IP range, the `192.168.1.1` and `192.168.1.2` addresses might already be in use. In this case, disconnect SSB from the network, and connect directly a computer to its external interface using a standard cross-link cable.

Step 7. To complete the procedure, click **Add**.

3.1.2. Procedure – Creating an alias IP address (Linux)

Purpose:

This procedure describes how to assign an alias IP address to a network interface on Linux platforms.

Steps:

Step 1. Start a terminal console (for example `gnome-terminal`, `konsole`, `xterm`, and so on).

Step 2. Issue the following command as root:

```
ifconfig <ethX>:0 192.168.1.2
```

where `<ethX>` is the ID of the network interface of the client, usually `eth0` or `eth1`.



- Step 3. Issue the `ifconfig` command. The `<ethX>:0` interface appears in the output, having `inet addr:192.168.1.2` .
- Step 4. Issue the `ping -c 3 192.168.1.1` command to verify that SSB is accessible. A similar result is displayed:

```
user@computer:~$ ping -c 3 192.168.1.1
PING 192.168.1.1 (192.168.1.1) 56(84) bytes of data.
64 bytes from 192.168.1.1: icmp-seq=1 ttl=63 time=0.357
ms
64 bytes from 192.168.1.1: icmp-seq=2 ttl=63 time=0.306
ms
64 bytes from 192.168.1.1: icmp-seq=3 ttl=63 time=0.314
ms
--- 192.168.1.1 ping statistics ---
3 packets transmitted, 3 received, 0% packet loss, time
2013ms
rtt min/avg/max/mdev = 0.306/0.325/0.357/0.030 ms
```

Open the page <https://192.168.1.1> from your browser and accept the certificate shown. The Welcome Wizard of SSB appears.

3.1.3. Procedure – Modifying the IP address of SSB

Purpose:

To configure SSB to listen for connections on a custom IP address, complete the following steps.



Warning

Use this procedure only before the initial configuration of SSB, that is, before completing the Welcome Wizard. For details on changing the IP address or other network settings of a configured SSB system, see *Section 4.3, Network settings* (p. 37).

Steps:

- Step 1. Access SSB from the local console.
- Step 2. Login using the `root` as username and `default` as password.
- Step 3. In the Console Menu, select **Shells > Core shell**.
- Step 4. Change the IP address of SSB using the following command. Replace `<IP-address>` with an IPv4 address suitable for your environment.
- ```
ifconfig eth0 <IP-address> netmask 255.255.255.0
```
- Step 5. Set the default gateway using the following command:
- ```
route add default gw <IP-address-of-the-default-gateway>
```
- Step 6. Type `exit`, then select **Logout** from the Console Menu.



Step 7. Open the page <https://<IP-address-you-set-for-SSB>> from your browser and accept the certificate shown. The Welcome Wizard of SSB appears.

3.2. Procedure – Configuring SSB with the Welcome Wizard

Purpose:

The Welcome Wizard guides you through the basic configuration steps of SSB. All parameters can be modified before the last step by using the **Back** button of the wizard, or later via the web interface of SSB.

Steps:

Step 1. Open the <https://<IP-address-of-SSB-external-interface>> page in your browser and accept the displayed certificate. The Welcome Wizard of SSB appears.



Tip

The SSB console displays the IP address the external interface is listening on. SSB either receives an IP address automatically via DHCP, or if a DHCP server is not available, listens on the 192.168.1.1 IP address.

Step 2. When configuring SSB for the first time, click **Next**.

Figure 3.6. The Welcome Wizard

It is also possible to import an existing configuration from a backup file. Use this feature to restore a backup configuration after a recovery, or to migrate an existing SSB configuration to a new device.

Step a. Click **Browse** and select the configuration file to import.



Note

It is not possible to directly import a GPG-encrypted configuration into SSB, it has to be decrypted locally first.



Step b. Enter the passphrase used when the configuration was exported into the **Encryption passphrase** field.

For details on restoring configuration from a configuration backup, see *Procedure 14.7, Restoring SSB configuration and data (p. 201)*

Step c. Click **Import**.



Warning

If you use the Import function to copy a configuration from one SSB to another, do not forget to configure the IP addresses of the second SSB. Having two devices with identical IP addresses on the same network leads to errors.

Step 3. Accept the End User License Agreement and install the SSB license

The screenshot shows the 'License' step of a 6-step installation wizard. The steps are: 1. Welcome, 2. License (current), 3. Networking, 4. Users, 5. Certificate, and 6. Finish.

License agreement

Software.

License Contract The present BalaBit Shell Control Box License Contract.

Product Any documentation referring to the BalaBit Shell Control Box or any module thereof, with special regard to the administration guide, the product description, the installation guide, user guides and manuals.

Documentation

Protected Host computers located in the zones protected by BalaBit Shell Control Box that have their administrative traffic controlled or audited by BalaBit Shell Control Box.

Servers

Protected The entire BalaBit Shell Control Box software including all of its modules, all the related Product Documentation; the source code, the structure of the databases, all registered

Confirmation

I have read and agree with the terms and conditions: ☒

License file upload

License file upload:

Customer: BalaBit IT Kft.

Serial: 308000-118577-92a06e

Host limit: Unlimited

Figure 3.7. The EULA and the license key

Step a. Read the End User License Agreement and select **Accept**.

Step b. Click **Browse**, select the SSB license file received with SSB, then click **Upload**.

Step c. Click **Next**.



Step 4. Fill the fields to configure networking. The meaning of each field is described below. The background of unfilled required fields is red. All parameters can later be modified using the regular interface of SSB.

Networking settings		
External interface:	IP address	Netmask
	10.50.0.111	255.255.255.0
Default GW:	10.50.0.254	
Hostname:	demo	
Domainname:	example.com	
DNS server:	10.50.0.254	
NTP server:	10.50.0.254	
SMTP server:	10.50.0.254	
Administrator's email:	documentation@example.com	
Timezone:	Europe/Budapest	
HA address:	Auto negotiation	

Back Next

Figure 3.8. Initial networking configuration

- Step a. **Hostname:** Name of the machine running SSB (for example SSB).
- Step b. **Domain name:** Name of the domain used on the network.
- Step c. **DNS server:** IP address of the name server used for domain name resolution.
- Step d. **NTP server:** The IP address or the hostname of the NTP server.
- Step e. **SMTP server:** The IP address or the hostname of the SMTP server used to deliver e-mails.
- Step f. **Administrator's e-mail:** E-mail address of the SSB administrator.
- Step g. **Timezone:** The timezone where the SSB is located.
- Step h. **External interface — IP address:** IP address of the external interface of SSB (for example 192.168.1.1). The IP address can be chosen from the range of the corresponding physical subnet. Clients will connect the external interface, therefore it must be accessible to them.

**Note**

Do not use IP addresses that fall into the following ranges:

- 1.2.0.0/16 (reserved for communication between SSB cluster nodes)
- 127.0.0.0/8 (localhost IP addresses)

Step i. **External interface — Netmask:** The IP netmask of the given range in IP format. For example, general class C networks have the 255.255.255.0 netmask.

Step j. **Default gateway:** IP address of the default gateway. When using several network cards, the default gateway is usually in the direction of the external interface.

Step k. **HA address:** The IP address of the high availability (HA) interface. Leave this field on *auto* unless specifically requested by the support team.

Step l. Click **Next**.

Step 5. Enter the passwords used to access SSB.

Figure 3.9. Passwords

**Note**

SSB passwords can contain the following special characters: !"#\$%&'()*+,-./:;<=>?@[\\]^_`{|}

Step a. **Admin password:** The password of the *admin* user who can access the web interface of SSB.

Step b. **Root password:** The password of the *root* user, required to access SSB via SSH or from the local console.

**Note**

Accessing SSB using SSH is rarely needed, and recommended only for advanced users for troubleshooting situations.

Step c. If you want to prevent users from accessing SSB remotely via SSH or changing the root password of SSB, select the **Seal the box** checkbox. Sealed mode can be activated later from the web interface as well. For details, see *Section 6.7, Sealed mode (p. 103)*.

Step d. Click **Next**.

Step 6. Upload or create a certificate for the SSB web interface. This SSL certificate will be displayed by SSB to authenticate administrative HTTPS connections to the web interface.

The screenshot shows the 'Certificate' step (5) of a 6-step process. The 'Server certificate' section shows a selected certificate for 'Server X.509 certificate' and 'Server private key'. The 'Generate new self-signed certificate' section has fields for Country (HU -- Hungary), Locality name (Budapest), Organization name (Example Inc), Organizational unit name (IT Security), and State or Province name. A 'Generate certificate' button is at the bottom of this section. 'Back' and 'Next' buttons are at the bottom of the form.

Figure 3.10. Creating a certificate for SSB

To create a self-signed certificate, fill the fields of the **Generate new self-signed certificate** section and click **Generate**. The certificate will be self-signed by the SSB appliance; the hostname of SSB will be used as the issuer and common name.

Step a. **Country**: Select the country where SSB is located (for example HU-Hungary).

Step b. **Locality**: The city where SSB is located (for example Budapest).


Step c. **Organization**: The company who owns SSB (for example Example Inc.).

Step d. **Organization unit**: The division of the company who owns SSB (for example IT Security Department).



Step e. **State or Province:** The state or province where SSB is located.

Step f. Click **Generate**.

If you want to use a certificate that is signed by an external Certificate Authority, in the **Server X.509 certificate** field, click  to upload the certificate.

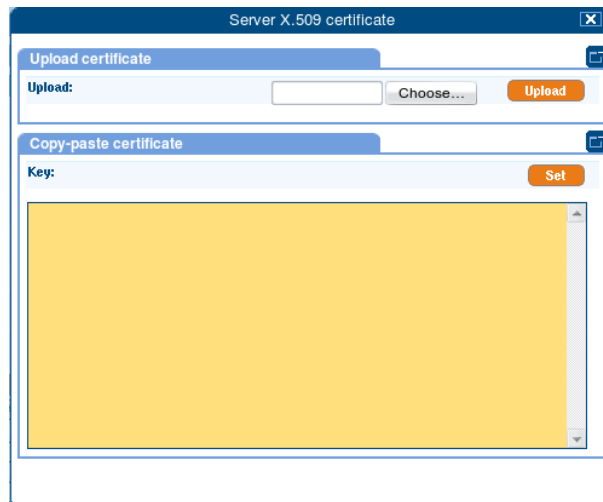



Figure 3.11. Uploading a certificate for SSB

Then in the **Server private key** field click , upload the private key, and enter the password protecting the private key.



Server private key Close

Upload key

Upload key: Upload

File: Browse...

Password:

Copy-paste key

Copy-paste key: Upload

Key:

Password:

Figure 3.12. Uploading a private key



Note

SSB accepts private keys in PEM (RSA and DSA), PUTTY, and SSHCOM/Tectia format. Password-protected private keys are also supported.

Step 7. Review the data entered in the previous steps. This page also displays the certificate generated in the last step; the RSA SSH key of SSB, and information about the license file.



Configuration details	
Mode:	bastion
Hostname:	demo
Domainname:	example.com
External address:	10.50.0.111
Management address:	10.50.0.111
Default gateway:	10.50.0.254
DNS server:	10.50.0.254
Syslog server:	10.50.0.2
Timezone:	Europe/Budapest
NTP server:	10.50.0.254
SSL certificate:	/C=HU/L=Budapest/O=Example Inc/OU=IT Security/CN=demo.example.com
SSH RSA key:	1024 57:c5:e6:18:fe:40:55:9c:9d:f5:82:19:c4:76:b2:8a
Licensed version:	2.0
Licensed customer:	BalaBit IT Kft.
Licensed hosts:	Unlimited
License serial:	308000-118577-92a06e

Figure 3.13. Review configuration data

If all information is correct, click **Finish**.

**Warning**

The configuration takes effect immediately after clicking **Finish**. Incorrect network configuration data can render SSB inaccessible.

SSB is now accessible from the regular web interface via the IP address of its external interface.

Step 8.

Figure 3.14. Logging in to SSB

Your browser is automatically redirected to the IP address set as the external interface of SSB, where you can login to the web interface of SSB using the *admin* username and the password you set for this user in the Welcome Wizard.



3.3. Procedure – Configuring storage access in SSB

Purpose:

The syslog-ng Store Box SANConnect edition of SSB uses a SAN storage module to store data. To configure SAN access in SSB, complete the following steps:

Steps:

Step 1. To configure the storage-related options of SSB, navigate to **Basic Settings > Storage**.



Note

If the Storage tab is not visible in the **Basic Settings** menu, the license file of SSB does not have the storage support enabled. If you have purchased an SSB version that includes storage support, contact your local distributor, or directly BalaBit S.a.r.l..

Step 2. *Optional step:* If you have SSB SANConnect version, and your SAN system has a name server, enter the IP address of the iSCSI Name Server into the ISNS field.

Step 3. Add an IP address and netmask (for example `1.2.4.135` and `255.255.255.0`) for the iSCSI card of the SSB. Make sure that the iSCSI interface of SSB and the storage module are on the same subnet. When running SSB in high availability mode, configure an IP address for the iSCSI card of the second SSB node as well. Make sure that the iSCSI interfaces of the two SSB nodes are different (for example use `1.2.4.136` and `255.255.255.0` for the second node).

Step 4. Add the storage module to the iSCSI Targets available from SSB. List every IP address and IQN of the storage module as an **iSCSI Target**. As a result, the volumes available on the storage will be displayed, including the LUN and WWN of every volume. The IP addresses and IQN numbers of the storage module are printed on the syslog-ng Store Box Certificate you received with the SSB hardware.



Note

In high availability mode, or if the iSCSI card of SSB has two ports connected to the storage module, every volume can be listed multiple times, and a volume may not be accessible from every target. This is normal.

Step 5. Format the volume if needed. Select **Format**, and enter a name for the volume. After the formatting procedure is finished, the size and the free space available on the volume are displayed.

Step 6. When creating logspaces, assign a volume to the logspace. Note that logspaces stored on a storage volume are identical to logspaces stored on the local hard drive: they can be shared, every logspace is stored in a separate directory, and so on. For details on creating logspaces, see *Section 8.4, Creating custom message spaces in SSB (p. 121)*.



Note

A single volume can store multiple logspaces, but a logspace can use only a single volume.




For details on managing SAN access from SSB, see *Section 6.11, Managing SAN access in SSB (p. 112)*.

For details on SAN troubleshooting, see *Procedure 14.8, SAN troubleshooting (p. 201)*.



Chapter 4. Basic settings

SSB is configured via the web interface. Configuration changes take effect automatically after clicking . Only the modifications of the current page or tab are activated — each page and tab must be committed separately.

- For details about the supported browsers, see *Section 4.1, Supported web browsers and operating systems (p. 33)*.
- For details on how to use the web interface of SSB, see *Section 4.2, The structure of the web interface (p. 33)*.
- For details on how to configure the basic settings of the SSB host like network settings, system monitoring, and backup settings, see *Section 4.3, Network settings (p. 37)*.

4.1. Supported web browsers and operating systems

Supported browsers: Mozilla Firefox 10 and Microsoft Internet Explorer 8 and 9. The browser must support HTTPS connections, JavaScript, and cookies. Make sure that both JavaScript and cookies are enabled.

Other tested browsers: Mozilla Firefox 3.6 and Google Chrome 17.

**Note**

SSB displays a warning message if your browser is not supported or JavaScript is disabled.

Supported operating systems: Microsoft Windows XP, Windows 2003 Server, Windows Vista, Windows 2008 Server, Windows 7, and Linux.


The SSB web interface can be accessed only using SSLv3 or TLSv1 encryption and strong cipher algorithms.

Note that when using Internet Explorer 7 on Windows 2008 to access SSB you must enable active scripting for the Internet Zone, otherwise the SSB web interface will not operate properly. This is not required on other platforms or browser versions. To accomplish this, select **Tools > Internet Options > Security > Internet Zone > Custom level**, and set the **Active scripting** field to *Enabled*.

When using Internet Explorer 8.0 on Windows Server 2008 R2 Enterprise, disable the Content Advisor in the Internet Explorer. To accomplish this, select **Tools > Internet Options > Content > Content Advisor > Disable**.

4.2. The structure of the web interface

The web interface consists of the following main sections:

Main menu: Each menu item displays its options in the main workspace on one or more tabs. Click  in front of a main menu item to display the list of available tabs.




Figure 4.1. Structure of the web interface

User menu: Provides possibilities to change your SSB password; to log out; and disable confirmation dialogs and tooltips using the **Preferences** option.

User info: Provides information about the user currently logged in:

- **User:** username
- **Host:** IP address of the user's computer
- **Last login:** date and IP address of the user's last login

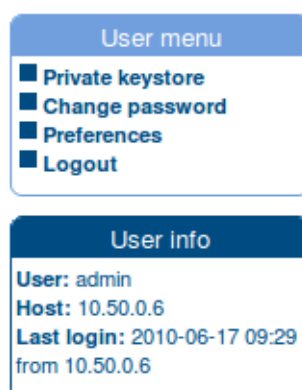


Figure 4.2. User menu and user info

System monitor: Displays accessibility and system health information about SSB, including the following:

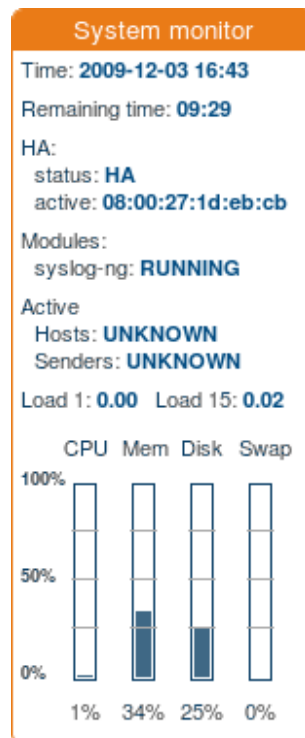


Figure 4.3. System monitor

- **Time:** System date and time.
- **Remaining time:** The time remaining before the session to the web interface times out.
- **Locked:** Indicates that the interface is locked by another administrator (for details, see *Section 4.2.2, Multiple web users and locking (p. 37)*)
- **Modules:** The status of syslog-ng running on SSB (ideally it is *RUNNING*).
- **License:** License information if the license is not valid, or an evaluation version license has expired.
- The status of the RAID devices, if synchronization between the disks is in progress.
- **Active:**
 - **Hosts:** the number of clients (log source hosts) where the log messages originate from (for example computers)
 - **Senders:** the number of senders where the log messages directly come from (for example relays)

**Example 4.1. Number of hosts and senders**

For example: if 300 clients all send log messages directly to SSB the Hosts and Senders are both 300.

If the 300 clients send the messages to 3 relays (assuming that the relays do not send messages themselves) and only the relays communicate directly with SSB then Hosts is 300, while Senders is 3 (the 3 relays).

If the relays also send messages, then Hosts is 303, while Senders is 3 (the 3 relays).



- **HA:** The HA status and the ID of the active node if two SSB units are running in a High Availability cluster. If there are redundant Heartbeat interfaces configured, their status is displayed as well. If the nodes of the cluster are synchronizing data between each other, the progress and the time remaining from the synchronization process is also displayed.
- Average system load during the
 - **Load 1:** last minute
 - **Load 15:** last fifteen minutes
- CPU, memory, hard disk, and swap use. Hover the mouse above the graphical bars to receive a more details in a tooltip, or navigate to **Basic Settings > Dashboard** for detailed reports.

The System monitor displays current information about the state of SSB. To display a history of these parameters, go to **Basic Settings > Dashboard**. For details, see *Section 14.5, Status history and statistics (p. 196)*.

4.2.1. Elements of the main workspace

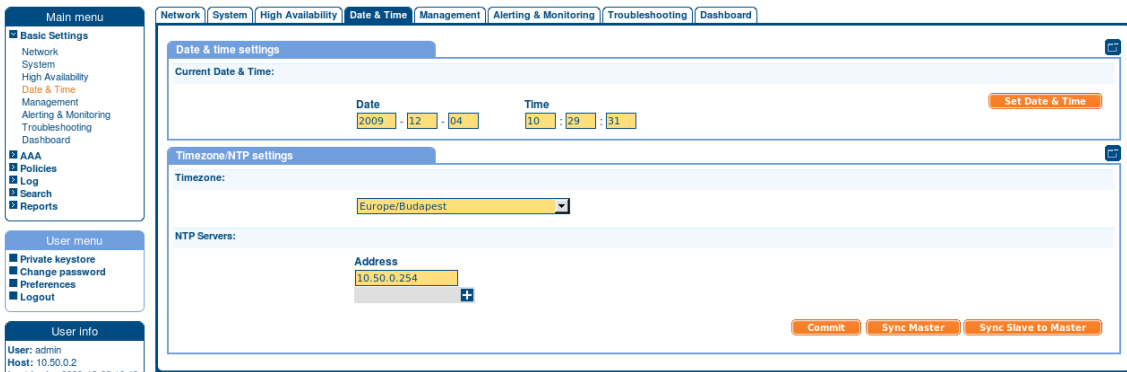







The main workspace displays the configuration settings related to the selected main menu item grouped into one or more tabs. Related parameters of a tab are organized into labeled groups or sections, marked with blue outline .

Figure 4.4. Main workspace

-  **Commit** Each page includes one or more orange action buttons. The most common action button is the **Commit**, which saves and activates the changes of the page.
-  **Show/Hide Details:** Displays or hides additional configuration settings and options.
-  **Create entry:** Create a new row or entry (for example an IP address or a policy).
-  **Delete entry:** Delete a row or an entry (for example an IP address or a policy).
-  **Open/collapse lists:** Open or close a list of options (for example the list of available reports).
-  **Modify entries or upload files:** Edit an entry (for example a host key, a list, and so on), or upload a file (for example a private key). These actions open a popup window where the actual modification can be performed.
-  **Position an item in a list:** Modify the order of items in a list. The order of items in a list (for example the order of connections, permitted channels in a channel policy, and so on) is important because when SSB is looking for a policy, it evaluates the list from top to down, and selects the first item completely matching the search criteria. For example, when a client initiates a connection to a



protected server, SSB selects the first connection policy matching the client's IP address, the server's IP address, and the target port (the From, To, and Port fields of the connection).

Message window: This popup window displays the responses of SSB to the user's actions, for example **Configuration saved successfully**. Error messages are also displayed here. All messages are included in the system log. For detailed system logs (including message history), see the **Troubleshooting** tab of the Basic menu. To make the window appear only for failed actions, navigate to **User menu > Preferences** and enable the **Autoclose successful commit messages** option.




Figure 4.5. Message window

4.2.2. Multiple web users and locking

Multiple administrators can access the SSB web interface simultaneously, but only one of them can modify the configuration. This means that the configuration of SSB is automatically locked when the first administrator who can modify the configuration accesses a configuration page (for example the **Basic Settings**, **AAA**, or **Logs** menu). The username and IP address of the administrator locking the configuration is displayed in the **System Monitor** field. Other administrators must wait until the locking administrator logs out, or the session of the administrator times out. However, it is possible to access the **Search** and **Reporting** menus, or browse the configuration with only View rights (for details, see *Section 5.6, Managing user rights and usergroups (p. 77)*).



Note

If an administrator logs in to SSB using the local console or a remote SSH connection, access via the web interface is completely blocked. Inactive local and SSH connections timeout just like web connections. For details, see *Section 6.4, Accessing the SSB console (p. 99)*.

4.3. Network settings

The **Network** tab contains the network interface and naming settings of SSB.



Main menu

- Basic Settings
 - Network
 - System
 - High Availability
 - Date & Time
 - Management
 - Alerting & Monitoring
 - Troubleshooting
 - Dashboard
- AAA
 - Policies
 - Log
 - Search
 - Reports

User menu

- Private key store
- Change password
- Preferences
- Logout

User info

User: admin
Host: 10.50.0.6
Last login: 2010-06-17 10:07 from 10.50.0.6

System monitor

Time: 2010-06-17 10:57
Remaining time: 07:38
Locked: admin@10.50.0.6
Modules: syslog-ng: STOPPED
Active Hosts: 0
Senders: 0
Load 1: 0 Load 15: 0
CPU Mem Disk Swap
100%
50%

Network settings

Interfaces

External interface:

Address	Netmask
10.50.0.200	255.255.255.0

Management enabled: ☒

Management interface:

Enable management interface: ☐

Address	Netmask

Routing table:

Address	Netmask	Gateway
0.0.0.0	0.0.0.0	10.50.0.254

Ethernet links

Interface	Speed	Link
Internal interface:	Auto negotiation	10Mb/s
External interface:	Auto negotiation	100Mb/s
Management interface:	Auto negotiation	10Mb/s
HA interface:		Link: <input checked="" type="checkbox"/>

Naming

Field	Value
Hostname:	ssbdemo
DNS search domain:	balabit
Primary DNS server:	10.10.0.1
Secondary DNS server:	

Figure 4.6. Network settings

- **External interface:** The Address and Netmask of the SSB network interface that receives client connections. Click the **+** and **x** icons to add new alias IP addresses (also called alias interfaces) or delete existing ones. At least one external interface must be configured. If the management interface is disabled, the SSB web interface can be accessed via the external interface. When multiple external interfaces are configured, the first one refers to the physical network interface, all others are alias interfaces. The SSB web interface can be accessed from all external interfaces (if no management interface is configured).

Optionally, you can enable access to the SSB web interface even if the management interface is configured by activating the **Management enabled** function.

**Warning**

If you enable management access on an interface and configure alias IP address(es) on the same interface, SSB will accept management connections only on the original address of the interface.

**Note**

Do not use IP addresses that fall into the following ranges:

- 1.2.0.0/16 (reserved for communication between SSB cluster nodes)
- 127.0.0.0/8 (localhost IP addresses)

**Note**

The speed of the interface is displayed for every interface. To explicitly set the speed of the interface, select the desired speed from the **Speed** field. Modifying the speed of the interface is recommended only for advanced users. Also note that changing the interface speed might not take effect if the network card of SSB has been replaced with one different from the original.

- **Management interface:** The Address and Netmask of the SSB network interface used to access the SSB web interface. If the management interface is configured, the web interface can be accessed only via this interface, unless access from other interfaces is explicitly enabled.

**Note**

Do not use IP addresses that fall into the following ranges:

- 1.2.0.0/16 (reserved for communication between SSB cluster nodes)
- 127.0.0.0/8 (localhost IP addresses)

4.3.1. Procedure – Configuring the management interface

Purpose:

To activate the interface, complete the following steps.

Steps:

Step 1. Navigate to **Basic Settings > Network > Interfaces**.

Figure 4.7. Configuring the management interface

Step 2. In the **Management interface** field, select **Enable management interface**.

Step 3. Into the **Address** field, enter the IP address of SSB's management interface.

Step 4. Into the **Netmask** field, enter the netmask related to the IP address.



Step 5.



Warning

After clicking **Commit**, the web interface will be available only via the management interface — it will not be accessible using the current (external) interface, unless the **Management enabled** option is selected for the external interface.

Ensure that the Ethernet cable is plugged and the management interface is connected to the network; this is indicated by a green check icon in the **Basic settings > Networks > Ethernet links > HA interface > Link** field. When using High Availability, ensure that the management interface of both SSB units is connected to the network.

The **HA interface** section indicates if a link is detected on the high availability interface.



Click **Commit**.

- **HA address:** The IP address of the high availability (HA) interface. Leave this field on *Auto negotiation* unless specifically requested by the support team.



Note

As of SSB version 1.1.1, when both nodes of a cluster boot up in parallel, the node with the **1.2.4.1** HA IP address will become the master node.

- **Interfaces > Routing table:** When sending a packet to a remote network, SSB consults the routing table to determine the path it should be sent. If there is no information in the routing table then the packet is sent to the default gateway. Use the routing table to define static routes to specific hosts or networks. You have to use the routing table if the internal interface is connected to multiple subnets, because the default gateway is (usually) towards the external interface. Click the  and  icons to add new routes or delete existing ones. A route means that messages sent to the **Address/Netmask** network should be delivered to **Gateway**.

For detailed examples, see *Procedure 4.3.2, Routing management traffic to the management interface (p. 40)*.

- **Naming > Hostname:** Name of the machine running SSB.
- **Naming > DNS search domain:** Name of the domain used on the network. When resolving the domain names of the audited connections, SSB will use this domain to resolve the target hostname if the appended domain entry of a target address is empty.
- **Naming > Primary DNS server:** IP address of the name server used for domain name resolution.
- **Naming > Secondary DNS server:** IP address of the name server used for domain name resolution if the primary server is unaccessible.

4.3.2. Procedure – Routing management traffic to the management interface

Purpose:

For security reasons — and also to reduce network usage on the external interface — it is recommended to direct all management-related traffic of SSB towards the management network interface. Such traffic includes



access to the web interface, *backups and archiving*, *data forwarded to a remote destination*, and *e-mail or SNMP alerts* sent to the administrator.

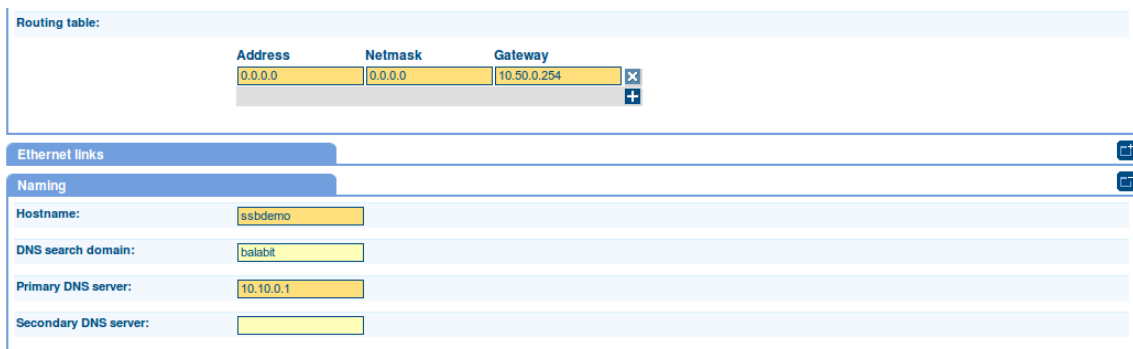


Warning

Complete the following procedure only if the management interface is configured; otherwise the data sent by SSB will be lost. For details on configuring the management interface, see *Procedure 4.3.1, Configuring the management interface* (p. 39).

Steps:

Step 1. To add a new routing entry, navigate to **Basic Settings > Network > Interfaces** and in the **Routing table** field, click .



Address	Netmask	Gateway
0.0.0.0	0.0.0.0	10.50.0.254

Ethernet links

Naming

Hostname: ssbdemo

DNS search domain: balabit

Primary DNS server: 10.10.0.1

Secondary DNS server:

Figure 4.8. Routing

- Step 2. Enter the IP address of the backup server (as set in *Procedure 4.7.1, Creating configuration and data backups* (p. 54)) into the **Address** field.
- Step 3. Enter the related netmask into the **Netmask** field.
- Step 4. Enter the IP address of the gateway used on that subnetwork into the **Gateway** field.
- Step 5. Click **Commit**.
- Step 6. Repeat Steps 1-5 and create a routing entry for other backup servers if needed.
- Step 7. Repeat Steps 1-5 and create a routing entry for the SMTP server (as set in *Section 4.5, SNMP and e-mail alerts* (p. 42)).
- Step 8. Repeat Steps 1-5 and create a routing entry for the remote destinations (as set in *Chapter 9, Forwarding messages from SSB* (p. 135)).

4.4. Date and time configuration

Date and time related settings of SSB can be configured on the **Date & Time** tab of the **Basic** page.



Figure 4.9. Date and time management



Warning

It is essential to set the date and time correctly on SSB, otherwise the date information of the logs will be inaccurate. SSB displays a warning on this page and sends an alert if the time becomes out of sync.

To explicitly set the date and time on SSB, enter the current date into respective fields of the **Date & Time Settings** group and click **Set Date & Time**.

4.4.1. Procedure – Configuring a time (NTP) server

Purpose:

To retrieve the date automatically from a time server, complete the following steps.

Steps:

- Step 1. Select your timezone in the **Timezone** field.
- Step 2. Enter the IP address of an NTP time server into the **Address** field.
- Step 3. Click **Commit**.
- Step 4. Click the **+** and **x** icons to add new servers or delete existing ones.



Note

If the time setting of SSB is very inaccurate (that is, the difference between the system time and the actual time is great), it might take a long time to retrieve the date from the NTP server. In this case, click **Sync now** to sync the time immediately using SNTP.

When two SSB units are operating in high availability mode, the **Sync now** button is named **Sync Master**, and synchronizes the time of the master node to the NTP server. To synchronize the time between the master and the slave nodes, click **Sync Slave to Master**.

4.5. SNMP and e-mail alerts

E-mail alerts can be configured on the **Basic Settings > Management** page.



The screenshot displays a web-based configuration interface for a network device. The top navigation bar includes tabs for Network, System, High Availability, Date & Time, Management, Alerting & Monitoring, Troubleshooting, and Dashboard. The left sidebar contains a 'Main menu' with options like Basic Settings, Network, System, High Availability, Date & Time, Management (highlighted), Alerting & Monitoring, Troubleshooting, and Dashboard. Below this is a 'User menu' with Private keystore, Change password, Preferences, and Logout. A 'User info' section shows the user is 'admin' with a last login on 2009-12-03 at 16:43. A 'System monitor' section displays system time, remaining time, HA status, active time, modules, active status, hosts, senders, and load.

The main configuration area is titled 'Management' and includes a 'Commit' button. It contains three sections:

- Syslog**: A section for configuring system logging.
- SNMP trap settings**: A section for configuring SNMP traps. It includes a text field for 'SNMP server address' (10.50.0.2), radio buttons for 'SNMP v2c' (selected) and 'SNMP v3', and a text field for 'Community' (public).
- SNMP agent settings**: A section for configuring SNMP agents. It includes text fields for 'Client address' (10.50.0.2), 'System location', 'System contact', and 'System description'. It also has checkboxes for 'SNMP v2c agent' (checked) and 'SNMP v3 agent' (unchecked), with a 'Community' field (public) for the v2c agent.
- Mail settings**: A section for configuring email alerts. It includes text fields for 'SMTP server address' (mail.example.com), 'Administrator's e-mail address' (admin@example.com), 'Send e-mail alerts to' (alert@example.com), and 'Send reports to' (boss@example.com). There is a 'Test' button next to the administrator's email address.

Figure 4.10. Configuring SNMP and e-mail alerts

4.5.1. Procedure – Configuring e-mail alerts

Purpose:

To configure e-mail alerts, complete the following steps:

Steps:

- Step 1. Navigate to **Basic Settings > Management > Mail settings**.
- Step 2. Enter the IP address or the hostname of the mail server into the **SMTP server address** field.



The screenshot shows the SSB web interface with the 'Alerting & Monitoring' tab selected. The 'Mail settings' section is expanded, displaying the following configuration fields:

- SMTP server address:** mail.example.com
- Administrator's e-mail address:** admin@example.com (with a 'Test' button next to it)
- Send e-mail alerts to:** alert@example.com
- Send reports to:** boss@example.com

The 'Commit' button is located in the top right corner of the configuration area. The left sidebar shows the 'Main menu' with 'Alerting & Monitoring' highlighted, and the 'User menu' with options like 'Private keystore', 'Change password', 'Preferences', and 'Logout'.

Figure 4.11. Configuring e-mail sending

- Step 3. Enter the e-mail address of the administrator into the **Administrator's e-mail address** field. SSB sends notifications related to system-events (but not alerts and reports) to this address.
- Step 4. Enter the e-mail address of the administrator into the **Send e-mail alerts to** field. SSB sends monitoring alerts to this address.
- Step 5. Enter the e-mail address the person who should receive traffic reports from SSB into the **Send reports to** field. For details on reports, see *Section 12.8, Reports (p. 178)*.



Warning

To get alert e-mails, provide an e-mail address in this field. Sending alerts fails if these settings are incorrect, since the alerting e-mail address does not fall back to the administrator's e-mail address by default.

- Step 6. Click **Commit**.
- Step 7. Click **Test** to send a test message.
If the test message does not arrive to the server, check if SSB can access the server. For details, see *Chapter 14, Troubleshooting SSB (p. 192)*.
- Step 8. Navigate to **Basic Settings > Alerting & Monitoring** and select in which situations should SSB send an e-mail alert. For details, see *Section 4.6, Configuring system monitoring on SSB (p. 47)*.
- Step 9. Click **Commit**.

4.5.2. Procedure – Configuring SNMP alerts

Purpose:

SSB can send alerts to a central monitoring server via SNMP (Simple Network Management Protocol). To configure SNMP alerts, complete the following steps:

Steps:

- Step 1. Navigate to **Basic Settings > Management > SNMP trap settings**.
- Step 2. Enter the IP address or the hostname of the SNMP server into the **SNMP server address** field.



Figure 4.12. Configuring SNMP alerts

Step 3. Select the SNMP protocol to use.

- To use the SNMP v2c protocol for SNMP queries, select **SNMP v2c**, and enter the community to use into the **Community** field.
- To use the SNMP v3 protocol, select **SNMP v3** and complete the following steps:

Figure 4.13. Configuring SNMP alerts using SNMPv3

Step a. Enter the username to use into the **Username** field.

Step b. Enter the engine ID to use into the **Engine ID** field. The engine ID is a hexadecimal number at least 10 digits long, starting with 0x. For example 0xABABABABAB.

Step c. Select the authentication method (**MD5 or SHA1**) to use from the **Authentication method** field.

Step d. Enter the password to use into the **Authentication password** field.

Step e. Select the encryption method (**Disabled, DES or AES**) to use from the **Encryption method** field.

Step f. Enter the encryption password to use into the **Encryption password** field.

Step 4. Click **Commit**.



Step 5. Navigate to **Basic Settings > Alerting & Monitoring** and select in which situations should SSB send an SNMP alert. For details, see *Section 4.6, Configuring system monitoring on SSB (p. 47)*.

Step 6. Click **Commit**.

4.5.3. Procedure – Querying SSB status information using agents

Purpose:

External SNMP agents can query the status information of SSB. To configure which clients can query this information, complete the following steps:

Steps:

Step 1. Navigate to **Basic Settings > Management > SNMP agent settings**.

Username	Auth. method	Auth. password	Encryption method	Encryption password
snmpagent	MD5	*****	AES	*****

Figure 4.14. Configuring SNMP agent access

Step 2. The status of SSB can be queried dynamically via SNMP. By default, the status can be queried from any host. To restrict access to these data to a single host, enter the IP address of the host into the **Client address** field.

Step 3. Optionally, you can enter the details of the SNMP server into the **System location**, **System contact**, and **System description** fields.

Step 4. Select the SNMP protocol to use.

- To use the SNMP v2c protocol for SNMP queries, select **SNMP v2c agent**, and enter the community to use into the **Community** field.
- To use the SNMP v3 protocol, select **SNMP v3 agent** and complete the following steps:

Step a. Click **+**

Step b. Enter the username used by the SNMP agent into the **Username** field.

Step c. Select the authentication method (**MD5 or SHA1**) to use from the **Auth. method** field.

Step d. Enter the password used by the SNMP agent into the **Auth. password** field.



Step e. Select the encryption method (**Disabled, DES or AES**) to use from the **Encryption method** field.

Step f. Enter the encryption password to use into the **Encryption password** field.

Step g. To add other agents, click **+**.

Step 5. Click **Commit**.

4.6. Configuring system monitoring on SSB

SSB continuously monitors a number of parameters of the SSB hardware and its environment. If a parameter reaches a critical level (set in its respective **Maximum** field), SSB sends e-mail and SNMP messages to alert the administrator.

SSB sends SNMP alerts using the management network interface by default, or using the external interface if the management interface is disabled. SSB supports the SNMPv2c and SNMPv3 protocols. The SNMP server set on the **Management** tab can query status information from SSB.



Tip

To have your central monitoring system recognize the SNMP alerts sent by SSB, select **Basic Settings > Alerting & Monitoring > Download MIBs** to download the SSB-specific Management Information Base (MIB), then import it into your monitoring system.



Main menu

- Basic Settings
 - Network
 - System
 - High Availability
 - Date & Time
 - Management
 - Alerting & Monitoring
 - Troubleshooting
 - Dashboard
- AAA
- Policies
- Log
- Search
- Reports

User menu

- Private keystore
- Change password
- Preferences
- Logout

User info

User: admin
Host: 10.40.255.254
Last login: 2011-10-19 13:12 from 10.40.255.254

System monitor

Time: 2011-10-19 14:01
Remaining time: 09:18
Locked: admin@10.40.255.254
Modules: syslog-ng: **RUNNING**
Active
Hosts: 2
Senders: 1
Load 1: 0.45 Load 15: 0.13
CPU Mem Disk Swap
100% 56% 9% 0%

Health monitoring

Disk utilization maximum: 80 %
Load 1 maximum: 5
Load 5 maximum: 4
Load 15 maximum: 3
Swap utilization maximum: 70 %

System related traps

Description	Name	Email	SNMP
Login failed	xcbLoginFailure	<input type="checkbox"/>	<input type="checkbox"/>
Successful login	xcbLogin	<input type="checkbox"/>	<input type="checkbox"/>
Logout from the management interface	xcbLogout	<input type="checkbox"/>	<input type="checkbox"/>
Configuration changed	xcbConfigChange	<input type="checkbox"/>	<input type="checkbox"/>
General alert	xcbAlert	<input type="checkbox"/>	<input type="checkbox"/>
General error	xcbError	<input type="checkbox"/>	<input type="checkbox"/>
Data and configuration backup failed	xcbBackupFailed	<input type="checkbox"/>	<input type="checkbox"/>
Data archiving failed	xcbArchiveFailed	<input type="checkbox"/>	<input type="checkbox"/>
Database error occurred	xcbDBError	<input type="checkbox"/>	<input type="checkbox"/>
Destination address limit reached	xcbLimitReached	<input type="checkbox"/>	<input type="checkbox"/>
HA node state changed	xcbHaNodeChanged	<input type="checkbox"/>	<input type="checkbox"/>
Timestamping error occurred	xcbTimestampError	<input type="checkbox"/>	<input type="checkbox"/>
Time sync lost	xcbTimeSyncLost	<input type="checkbox"/>	<input type="checkbox"/>
Raid status changed	xcbRaidStatus	<input type="checkbox"/>	<input type="checkbox"/>
Hardware error occurred	xcbHWEError	<input type="checkbox"/>	<input type="checkbox"/>
Firmware is tainted	xcbFirmwareTainted	<input type="checkbox"/>	<input type="checkbox"/>
Disk usage is above the defined ratio	xcbDiskFull	<input type="checkbox"/>	<input type="checkbox"/>

syslog-ng traps

Description	Name	Email	SNMP
syslog-ng failure	syslogngFailureTrap	<input type="checkbox"/>	<input type="checkbox"/>
Remote syslog-ng peer configuration changed	peerConfigChangeTrap	<input type="checkbox"/>	<input type="checkbox"/>
Logspace exceeded warning size	spaceSizeLimit	<input type="checkbox"/>	<input type="checkbox"/>
Message rate was outside the specified limits	ssbAbsoluteMessageRateAlert	<input type="checkbox"/>	<input type="checkbox"/>
Too many message rate alerts were generated	ssbRateLimitTooManyAlerts	<input type="checkbox"/>	<input type="checkbox"/>

Figure 4.15. Configuring SNMP and e-mail alerting

4.6.1. Procedure – Configuring monitoring

Purpose:

To configure monitoring, complete the following steps:

Steps:

- Step 1. Navigate to **Basic Settings > Alerting & Monitoring**.
- Step 2. The default threshold values of the parameters are suitable for most situations. Adjust the thresholds only if needed.
- Step 3. Click **Commit**.
- Step 4. Navigate to **Basic Settings > Management** and verify that the **SNMP settings** and **Mail settings** of SSB are correct. SSB sends alerts only to the alert e-mail address and to the SNMP server.

**Warning**

Sending alerts fails if these settings are incorrect.

The following sections describe the parameters you can receive alerts on.

- For details on health-monitoring alerts, see *Section 4.6.2, Health monitoring (p. 49)*.
- For details on system-monitoring alerts, see *Section 4.6.5, System related traps (p. 52)*.
- For details on syslog-related alerts, see *Section 4.6.6, Alerts related to syslog-ng (p. 53)*.

4.6.2. Health monitoring

- **Disk utilization maximum:** Ratio of free space available on the hard disk. SSB sends an alert if the log files use more space than the set value. Archive the log files to a backup server to free disk space. For details, see *Procedure 4.7.2, Archiving the collected data (p. 59)*.

**Note**

The alert message includes the actual disk usage, not the limit set on the web interface. For example, you set SSB to alert if the disk usage increases above 10 percent. If the disk usage of SSB increases above this limit (for example to 17 percent), you receive the following alert message: *less than 90% free (= 17%)*. This means that the amount of used disk space increased above 10% (what you set as a limit, so it is less than 90%), namely to 17%.

- **Load 1|5|15 maximum:** The average load of SSB during the last one, five, or 15 minutes.
- **Swap utilization maximum:** Ratio of the swap space used by SSB. SSB sends an alert if it uses more swap space than the set value.

4.6.3. Procedure – Preventing disk space fill up

Purpose:

To prevent disk space from filling up, complete the following steps:

Steps:

- Step 1. Navigate to **Basic Settings > Management > Disk space fill up prevention**.
- Step 2. Set the limit of maximum disk utilization in percents in the respective field. When disk space is used above the set limit, SSB disconnects all clients. Entering 0 turns the feature off. The default value is 0.
- Step 3. *Optional step:* Enable the **Automatically start archiving** option to automatically start all configured archiving/cleanup jobs when disk usage goes over the limit.

**Note**

If there is no archiving policy set, enabling this option will not trigger automatic archiving.

Step 4. Navigate to **Basic Settings > Alerting & Monitoring > System related traps** and enable alert **Disk usage is above the defined ratio**.

Step 5. Click **Commit**.

4.6.4. Procedure – Configuring message rate alerting

Purpose:

With message rate alerting, you can detect the following abnormalities in SSB:

- The syslog-ng inside SSB has stopped working.
- One of the clients/sites sending logs is not detectable.
- One of the clients/sites is sending too many logs, probably unnecessarily.

Message rate alerting can be set for sources, spaces and destinations (remote or local).

Steps:

Step 1. Navigate to **Log** and select **Sources**, **Spaces** or **Destinations**.

Step 2. Enable **Message rate alerting**.

Step 3. In case of **Sources**, select the counter to be measured:

- *Messages*: Number of messages
- *Messages/sender*: Number of messages per sender (the last hop)
- *Messages/hostname*: Number of messages per host (based on the hostname in the message)

In case of **Spaces** or **Destinations**, the counter is the number of messages.

Step 4. Select the time period (between 5 minutes and 24 hours) during which the range is to be measured.

Step 5. Enter the range that is considered normal in the **Minimum** and **Maximum** fields.

Step 6. Select the alerting frequency in the **Alert** field. **Once** sends only one alert (and after the problem is fixed, a "Fixed" message), **Always** sends an alert each time the result of the measurement falls outside the preset range.

**Example 4.2. Creating an early time alert**

In case you want an early time alert, can create a normal (non master) alert with a very low minimum number of messages and a low check interval.



Counter	Period	Minimum	Maximum	Alert	Master alert
Messages	24 hours	10000	1000000	Once	<input type="checkbox"/>
Messages	30 minutes	10	1000	Once	<input type="checkbox"/>

Figure 4.16. Creating an early time alert

Step 7. If you have set more than one message rate alerts, you can set a master alert where applicable. To set an alert to be a master alert, select the **Master alert** checkbox next to it.

When a master alert is triggered (and while it remains triggered), all other alerts for the given source/destination/space are suppressed. A master alert only blocks the other alerts that would be triggered at the given timeslot. A 24-hour alert does not block alerts that would be triggered at, for example 00:05.

Suggestions for setting the master alert:

- set the master alert to low check interval (5 minutes, if possible)
- set the master alert to a lower check interval than the alerts it suppresses
- set the master alert to have more lax limits than the alerts it suppresses

The following examples demonstrate a few common use cases of a **Master alert**.



Example 4.3. Using the master alert to indicate unexpected events

The user has 2 relays (sender) and 10 hosts per each relay (=20 hosts). Each host sends approximately 5-10 messages in 5 minutes. Two message rate alerts are set, and one master alert to signal extreme unexpected events. Such event can be that either a host is undetectable and probably has stopped working, or that it sends too many logs, probably due to an error. The following configuration helps detecting these errors without having to receive hundreds of alerts unnecessarily.

Counter	Period	Minimum	Maximum	Alert	Master alert
Messages/hostname	5 minutes	50	100	Once	<input type="checkbox"/>
Messages/sender	5 minutes	500	1000	Once	<input type="checkbox"/>
Messages	5 minutes	0	10000	Once	<input checked="" type="checkbox"/>

Figure 4.17. Using the master alert to indicate unexpected events

Step 8. *Optional step:* Global alerts count the number of all messages received by syslog-ng on all sources, including internal messages.

- Step a. Navigate to **Log > Options > Message rate alerting statistics**. To add a global alert, click **+** at **Global alerts**.
- Step b. Select the time period (between 5 minutes and 24 hours) during which the range is to be measured.
- Step c. Enter the range that is considered normal in the **Minimum** and **Maximum** fields.
- Step d. Select the alerting frequency in the **Alert** field. **Once** sends only one alert (and after the problem is fixed, a "Fixed" message), **Always** sends an alert each time the result of the measurement falls outside the preset range.



Step e. To set the alert as a system-wide master alert, select **Global master alert**. It will suppress all other log rate alerts on SSB when it is triggered.



Note

In the following cases, a so-called "always"-type super-master alert is triggered automatically.

If all or some of the statistics from syslog-ng cannot be fetched, an alert is sent out and all other errors are suppressed until the error is fixed.

If, for some reason, syslog-ng sends an unprocessable amount of statistics (for example because of some invalid input data), a similar super-master alert is triggered and stops processing the input.

Step 9. *Optional step*: Navigate to **Log > Options > Message rate alerting statistics**. Set the maximum number of alerts you want to receive in **Limit of alerts sent out in a batch** to prevent alert flooding. SSB will send alerts up to the predefined value and then one single alert stating that too many message alerts were generated and the excess amount have not been sent.



Warning

Hazard of data loss! The alerts over the predefined limit will be unreachable.

4.6.5. System related traps

Name	SNMP alert ID	Description
Login failed	<i>xcbLoginFailure</i>	Failed login attempts from SSB web interface.
Successful login	<i>xcbLogin</i>	Successful login attempts into SSB web interface.
Logout from the management interface	<i>xcbLogout</i>	Logouts from SSB web interface.
Configuration changed	<i>xcbConfigChange</i>	Any modification of SSB's configuration.
General alert	<i>xcbAlert</i>	General alerts and error messages occurring on SSB. Note, that alerts on general alerts and errors are sent whenever there is an alert or error level message in the SSB system log. These messages are very verbose and mainly useful only for debugging purposes.
General error	<i>xcbError</i>	



Name	SNMP alert ID	Description
		Enabling these alerts may result in multiple e-mails or SNMP traps sent about the same event.
Data and configuration backup failed	<i>xcbBackupFailed</i>	Alerts if the backup procedure is unsuccessful.
Data archiving failed	<i>xcbArchiveFailed</i>	Alerts if the archiving procedure is unsuccessful.
Database error occurred	<i>xcbDBError</i>	An error occurred in the database where SSB stores alerts and accounting information. Contact our support team (see <i>Section 5, Contact and support information (p. xiii)</i> for contact information).
License limit reached	<i>xcbLimitReached</i>	Maximum number of clients has been reached.
HA node state changed	<i>xcbHaNodeChanged</i>	A node of the SSB cluster changed its state, for example, a takeover occurred.
Timestamping error occurred	<i>xcbTimestampError</i>	An error occurred during the timestamping process, for example the timestamping server did not respond.
Time sync lost	<i>xcbTimeSyncLost</i>	The system time became out of sync.
Raid status changed	<i>xcbRaidStatus</i>	The status of the node's RAID device changed its state.
Hardware error occurred	<i>xcbHWEError</i>	SSB detected a hardware error.
Firmware is tainted	<i>xcbFirmwareTainted</i>	A user has locally modified a file from the console.
Disk usage is above the defined ratio	<i>xcbDiskFull</i>	Disk space is used above the limit set in Disk space fill up prevention .

Table 4.1. System related traps

4.6.6. Alerts related to syslog-ng

Name	SNMP alert ID	Description
syslog-ng failure	<i>syslogngFailureTrap</i>	The syslog-ng application did not start properly, shut down unexpectedly, or encountered another problem. Depending on the error, SSB may not accept incoming



Name	SNMP alert ID	Description
		messages or send them to the destinations.
Remote syslog-ng peer configuration changed	<i>peerConfigChangeTrap</i>	The configuration of the syslog-ng application running on a remote host that sends its logs to SSB has been changed. Note that such changes are detected only if the remote peer uses at least version 3.0 of syslog-ng or version 3.0 of the syslog-ng Agent, and if messages from the <i>internal</i> source are sent to SSB.
Logspace exceeded warning size	<i>spaceSizeLimit</i>	The size of a log space has exceeded the size set as warning limit.

Table 4.2. Alerts related to syslog-ng

4.7. Data and configuration archiving and backups

The syslog-ng Store Box can create automatic backups of its configuration and the stored logs to a remote server. Logs can be archived as well. Backups and archiving is controlled using backup and archiving policies that define the protocol to use, the address of the backup server, and other parameters.

4.7.1. Procedure – Creating configuration and data backups

Purpose:

To configure automatic configuration backups, complete the following steps:



The screenshot displays the Mikrotik WinBox interface for configuring backup policies. On the left, the 'Main menu' is visible with options like Basic Settings, AAA, Policies, Hostlists, Shares, Backup & Archive/Cleanup, Log, Search, and Reports. Below it is the 'User menu' with Private keystore, Change password, Preferences, and Logout. The 'User info' section shows the user is 'admin' with host '10.50.0.2' and a last login on '2009-12-04 10:25'. The 'System monitor' section shows the time as '2009-12-03 16:43', remaining time as '09:29', HA status as 'HA', and modules as 'syslog-ng: RUNNING'. The main content area is titled 'Backup & Archive/Cleanup' and contains the 'Backup policies' section. A policy named 'main-backup' is being configured. The 'Start time' is set to '23:30'. Under 'Target settings', 'NFS' is selected as the protocol, with 'Target server' set to 'backup.example.co' and 'Export' set to 'backups'. Under 'Send notification', 'Send notification on all events' is selected. There are 'Commit' buttons at the top right and bottom right of the configuration area.

Figure 4.18. Configuring backups

Steps:

Step 1. Create a backup policy.

- Step a. Navigate to **Policies > Backup & Archive/Cleanup** and click **+** in the **Backup policies** section to create a new backup policy.
- Step b. Enter a name for the backup policy (for example *config-backup*).
- Step c. Enter the time when the backup process should start into the **Start time** field in HH:MM format (for example *23:00*).
- Step d. Enter the IP address or the hostname of the remote server into the **Target server** field (for example *backup.example.com*).



Step e.

Figure 4.19. Configuring backups

SSB can access the remote server via different protocols. Select the one to use from the available protocols:

- **Rsync over SSH:** Execute the rsync command via the Secure Shell protocol. Note that the backup server must run rsync version 3.0 or newer.
- **SMB/CIFS:** Server Message Block protocol used on Microsoft Windows Network.

**Warning**

The CIFS implementation of NetApp storage devices is not compatible with the CIFS implementation used in SSB, therefore it is not possible to create backups and archives from SSB to NetApp devices using the CIFS protocol (the operation fails with a similar error message: `/opt/scb/mnt/14719217504d41370514043/reports/2010": Permission denied (13) '2010/day/' rsync: failed to set times on`).

To overcome this problem, either:

- use the NFS protocol to access your NetApp devices, or
- use a backup device that has a CIFS implementation compatible with SSB, for example, Windows or Linux Samba.

**Warning**

When using the CIFS protocol to backup or archive files to a target server running Windows 2008 R2 that uses NTLMv2 authentication, the operation may fail with a similar error message:



```
CIFS VFS: Unexpected SMB signature
Status code returned 0xc000000d
NT_STATUS_INVALID_PARAMETER
CIFS VFS: Send error in SessSetup = -22
CIFS VFS: cifs_mount failed w/return code = -22
CIFS VFS: Server requires packet signing to be enabled
in /proc/fs/cifs/SecurityFlags.
CIFS VFS: cifs_mount failed w/return code = -95
CIFS VFS: Server requires packet signing to be enabled
in /proc/fs/cifs/SecurityFlags.
CIFS VFS: cifs_mount failed w/return code = -95
```

To overcome this problem, either:

- use the NFS protocol to access your Windows 2008 R2 servers, or
- edit the registry of the Windows 2008 R2 server or apply a hotfix. For details, see [Article 957441](#) in the Microsoft® Support site.

■ NFS: Network File System protocol.



Warning

When using the NFS protocol to create backups or archives, ensure that the files on the remote server are readable for the *www-data* user as well, because SSB uses this user to access remote backups and archives if needed.

Step f. Provide the protocol-specific parameters for the selected method. The protocol-specific parameters are described in *Section 4.7.4, Parameters of the backup protocols (p. 63)*.

Step g. To receive e-mail notification of the backup, select the **Send notification on errors only** or the **Send notification on all events** option. Notifications are sent to the administrator e-mail address set on the **Management** tab, and include the list of the files that were backed up.



Warning

Starting with SSB 3.0, the notification e-mail does not include the list of backed up file. To include the list of files in the e-mail, select **Send notification on all events** and disable the **Omit file list** option. However, note that if list is very long (for example, SSB stores over 20000 audit trails), the SSB web interface might become inaccessible.



Note

This e-mail notification is different from the one set on the **Alerting & Monitoring** tab. This notification is sent to the administrator's e-mail address, while the alerts are sent to the alert e-mail address (see *Section 4.6, Configuring system monitoring on SSB (p. 47)*).

Step h. Click **Commit**.

Expected outcome:

A backup policy is created.



Step 2. To use this policy to create configuration backups, navigate to **Basic Settings > Management > System backup**, and select the backup policy you want to use for backing up the configuration of SSB in the **System backup policy** field.

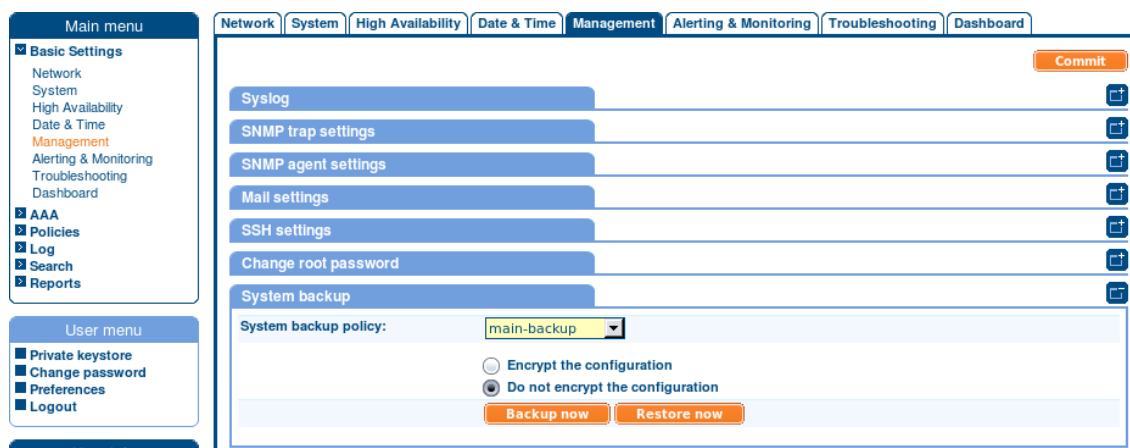


Figure 4.20. Configuring system backups

To use this policy to create data backups, navigate to **Log > Spaces**, select the logspace you want to backup, and select a backup policy in the **Backup policy** field.

Step 3. Click **Commit**.



Tip

To create an immediate backup of SSB's configuration to your machine (not to the backup server), select **Basic Settings > System > Export configuration**.

Note that the configuration export contains only the system settings and configuration files (including changelogs). System backups includes additional information like reports and alerts.

For details on restoring configuration from a configuration backup, see *Procedure 14.7, Restoring SSB configuration and data* (p. 201)



Note

Backup is different from archiving: the purpose of backup is to create a snapshot of SSB's configuration or the data stored on SSB that can be used for recovery in case of errors. Backup deletes all other data from the target directory; while restoring a backup deletes all other data from SSB.



Tip

To start the backup process immediately, click **Backup now**. The **Backup now** functionality works only after a backup policy has been selected.

To restore the stored data (logs, reports, and so on), click **Restore now**. Note that the **Restore now** function does not restore the configuration files of SSB.

When restoring a data backup, you must also import the SSB configuration relevant to the time when the backup was created (restoring the complete system backup is not necessary). Otherwise, SSB will not handle the restored data properly.

**Warning**

Before restoring data, make sure to stop the incoming log traffic, otherwise the log files of the current day might become inconsistent, preventing SSB from accepting messages for logspaces that had already received messages on the current day.

- To stop receiving log messages completely, select **Basic Settings > System > System Control > Syslog traffic > Disable**.
- To stop sending messages only to selected logspaces, select **Log > Paths**, and uncheck the **Enabled** field for the particular logspaces.

4.7.2. Procedure – Archiving the collected data

Purpose:

To configure data archiving, complete the following steps.

Steps:

Step 1. Create an archive policy.

Step a. Navigate to **Policies > Backup & Archive/Cleanup** and click **+** in the **Archive/Cleanup policies** section to create a new archive policy.

Step b.

Figure 4.21. Configuring backups and archiving

Enter a name for the archive policy.

Step c. Enter the time when the archive process should start into the **Start time** field in HH:MM format (for example 23:00).

Step d. SSB can access the remote server via different protocols. Select the one to use from the available protocols:



- **Only cleanup, no archiving:** Do not archive data to a server, simply delete the data that is older than **Retention time in days**.



Warning

No archiving permanently deletes all log files and data that is older than **Retention time in days** without creating a backup copy or an archive. Such data is irrecoverably lost. Use this option with care.

- **SMB/CIFS:** Server Message Block protocol used on Microsoft Windows Network.



Warning

The CIFS implementation of NetApp storage devices is not compatible with the CIFS implementation used in SSB, therefore it is not possible to create backups and archives from SSB to NetApp devices using the CIFS protocol (the operation fails with a similar error message: `/opt/scb/mnt/14719217504d41370514043/reports/2010": Permission denied (13) '2010/day/' rsync: failed to set times on`).

To overcome this problem, either:

- use the NFS protocol to access your NetApp devices, or
- use a backup device that has a CIFS implementation compatible with SSB, for example, Windows or Linux Samba.



Warning

When using the CIFS protocol to backup or archive files to a target server running Windows 2008 R2 that uses NTLMv2 authentication, the operation may fail with a similar error message:

```
CIFS VFS: Unexpected SMB signature
Status code returned 0xc000000d
NT_STATUS_INVALID_PARAMETER
CIFS VFS: Send error in SessSetup = -22
CIFS VFS: cifs_mount failed w/return code = -22
CIFS VFS: Server requires packet signing to be enabled
in /proc/fs/cifs/SecurityFlags.
CIFS VFS: cifs_mount failed w/return code = -95
CIFS VFS: Server requires packet signing to be enabled
in /proc/fs/cifs/SecurityFlags.
CIFS VFS: cifs_mount failed w/return code = -95
```

To overcome this problem, either:

- use the NFS protocol to access your Windows 2008 R2 servers, or
- edit the registry of the Windows 2008 R2 server or apply a hotfix. For details, see [Article 957441](#) in the Microsoft® Support site.

- **NFS:** Network File System protocol.

**Warning**

When using the NFS protocol to create backups or archives, ensure that the files on the remote server are readable for the *www-data* user as well, because SSB uses this user to access remote backups and archives if needed.

Step e. Enter the IP address or the hostname of the remote server into the **Target settings** > **Target server** field (for example *backup.example.com*).

Step f. Fill the **Retention time in days** field. Data older than this value is archived to the external server.

**Note**

The archived data is deleted from SSB.

Step g. Provide the protocol-specific parameters for the selected method. The protocol-specific parameters are described in *Section 4.7.4, Parameters of the backup protocols (p. 63)*.

Step h. To receive e-mail notification of the backup, select the **Send notification on errors only** or the **Send notification on all events** option. Notifications are sent to the administrator e-mail address set on the **Management** tab, and include the list of the files that were backed up.

**Note**

This e-mail notification is different from the one set on the **Alerting & Monitoring** tab. This notification is sent to the administrator's e-mail address, while the alerts are sent to the alert e-mail address (see *Section 4.6, Configuring system monitoring on SSB (p. 47)*).

Step i. Click **Commit**.

Expected outcome:

An archive policy is created.

Step 2. To use this policy to archive the data of a logspace, navigate to **Log** > **Spaces**, select the logspace you want to archive, and select the archive policy you want to use in the **Archive/Cleanup policy** field.

**Note**

If you change the connection protocol used for archiving (for example from NFS to SMB/CIFS), the old archives will become inaccessible. To avoid this, create a new archive policy using the new connection protocol, and select it at all affected connections (<**connection type**> > **Connections** > >**Archive/Cleanup policy**). This way, both the old and the new archived trails will be accessible.



Step 3. Click **Commit**.

**Tip**

To start the archiving process immediately, click **Archive now**. The **Archive now** functionality works only after the archiving has been configured.

4.7.3. Procedure – Encrypting configuration backups with GPG

Purpose:

To encrypt the configuration file of SSB during system-backups using the public-part of a GPG key. To enable GPG-encryption, complete the following steps:

**Warning**

The system-backups of SSB contain other information as well (for example databases), but only the configuration file is encrypted. Note that system-backups do not contain data like logspaces .

Steps:

Step 1. Navigate to **Basic > System > Management > System backup**.

Step 2. Select **Encrypt configuration**.

Step 3. Select ☒.

- To upload a key file, click **Browse**, select the file containing the public GPG key, and click **Upload**. SSB accepts both binary and ASCII-armored GPG keys.
- To copy-paste the key from the clipboard, paste it into the **Key** field and click **Set**.

**Note**

The GPG key you upload must be permitted to encrypt data. Keys that can be used only for signing cannot be used to encrypt the configuration file.

Step 4. Click **Commit**.

**Note**

It is not possible to directly import a GPG-encrypted configuration into SSB, it has to be decrypted locally first.

For details on restoring configuration from a configuration backup, see *Procedure 14.7, Restoring SSB configuration and data* (p. 201)



4.7.4. Parameters of the backup protocols

This section describes the details of the protocols used for data backup and archiving.

- For details on using Rsync, see *Procedure 4.7.4.1, Configuring Rsync over SSH* (p. 63).
- For details on using the Samba protocol, see *Procedure 4.7.4.2, Configuring SMB* (p. 64).
- For details on using NFS, see *Procedure 4.7.4.3, Configuring NFS* (p. 66).

4.7.4.1. Procedure – Configuring Rsync over SSH

Purpose:

The **Rsync over SSH** backup method connects the target server with SSH and executes the rsync UNIX command to copy the data to the remote server. SSB authenticates itself with a public key — password-based authentication is not supported. To configure this method, complete the following steps.



Warning

The backup server must run rsync version 3.0 or newer.

Steps:

Step 1. Select **Rsync over SSH** from the **Target settings** radio buttons.

The screenshot shows the 'Backup & Archive/Cleanup' configuration page. On the left, there are three sidebars: 'Main menu' with links like Basic Settings, AAA, Policies, Hostlists, Shares, Backup & Archive/Cleanup, Log, Search, and Reports; 'User menu' with Private keystore, Change password, Preferences, and Logout; and 'User info' showing user details. The main area is titled 'Backup policies' and contains two tabs: 'main-backup' and 'rsync-backup'. The 'rsync-backup' tab is active, showing 'Start time' set to 22:00. Under 'Target settings', 'Rsync over SSH' is selected with radio buttons. Fields include 'Username' (backup), 'Target server' (backup.example.com), 'Authentication key' (a long hex string), 'Server host key' (another long hex string), and 'Path' (/backups/). At the bottom, 'Send notification' options are shown, with 'Send notification on all events' selected.

Figure 4.22. Configuring backups using rsync

Step 2. Enter the username used to logon to the remote server into the **Username** field.

Step 3. Click ☒ in the **Authentication key** field. A popup window is displayed.



- Step 4. Generate a new keypair by clicking **Generate** or upload or paste an existing one. This key will be used to authenticate SSB on the remote server. The public key of this keypair must be imported to the remote server.
- Step 5. Click ☒ in the **Server host key** field. A popup window is displayed.
- Step 6.

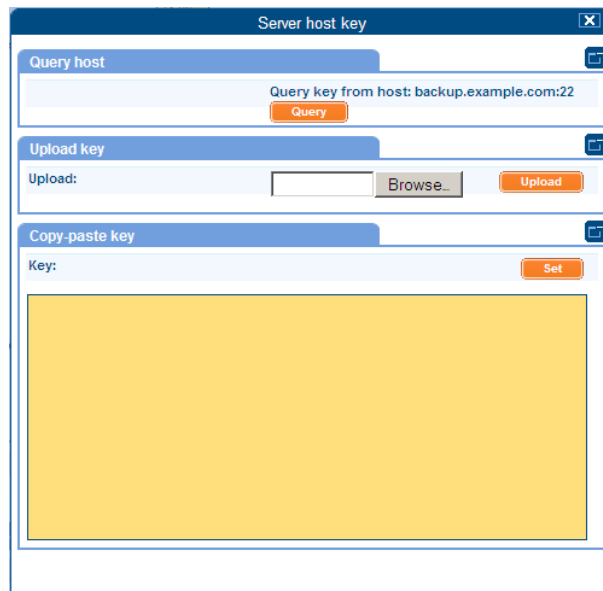


Figure 4.23. Configuring SSH keys

Click **Query** to download the host key of the server, or upload or paste the host key manually. SSB will compare the host key shown by the server to this key, and connect only if the two keys are identical.

- Step 7. Enter the port number of the SSH server running on the remote machine into the **Port** field.
- Step 8. Enter the path to the backup directory on the target server into the **Path** field (for example `/backups`). SSB saves all data into this directory, automatically creating subdirectories for logspaces. As a result of this, the same backup policy can be used for multiple logspaces. To ensure that a restore can be performed even if the logspace has been renamed, the subdirectories are created using a persistent internal ID of the logspace. To facilitate manual debugging, a text file is also saved in the directory with the name of the logspace, containing the internal ID for the logspace. This text file is only provided for troubleshooting purposes and is not used by SSB in any way.
- Step 9. Click **Commit**.

4.7.4.2. Procedure – Configuring SMB

Purpose:

The **SMB/CIFS** backup method connects to a share on the target server with Server Message Block protocol. SMB/CIFS is mainly used on Microsoft Windows Networks. To configure this method, complete the following steps.

**Warning**

The CIFS implementation of NetApp storage devices is not compatible with the CIFS implementation used in SSB, therefore it is not possible to create backups and archives from SSB to NetApp devices using the CIFS protocol (the operation fails with a similar error message: `/opt/scb/mnt/14719217504d41370514043/reports/2010": Permission denied (13) '2010/day/' rsync: failed to set times on).`

To overcome this problem, either:

- use the NFS protocol to access your NetApp devices, or
- use a backup device that has a CIFS implementation compatible with SSB, for example, Windows or Linux Samba.

**Warning**

When using the CIFS protocol to backup or archive files to a target server running Windows 2008 R2 that uses NTLMv2 authentication, the operation may fail with a similar error message:

```
CIFS VFS: Unexpected SMB signature
Status code returned 0xc000000d NT_STATUS_INVALID_PARAMETER
CIFS VFS: Send error in SessSetup = -22
CIFS VFS: cifs_mount failed w/return code = -22
CIFS VFS: Server requires packet signing to be enabled in /proc/fs/cifs/SecurityFlags.
CIFS VFS: cifs_mount failed w/return code = -95
CIFS VFS: Server requires packet signing to be enabled in /proc/fs/cifs/SecurityFlags.
CIFS VFS: cifs_mount failed w/return code = -95
```

To overcome this problem, either:

- use the NFS protocol to access your Windows 2008 R2 servers, or
- edit the registry of the Windows 2008 R2 server or apply a hotfix. For details, see [Article 957441](#) in the Microsoft® Support site.

Steps:

Step 1. Select **SMB/CIFS** from the **Target settings** radio buttons.



Figure 4.24. Configuring backups via SMB/CIFS

- Step 2. Enter the username used to logon to the remote server into the **Username** field, or select the **Anonymous** option.
- Step 3. Enter the password corresponding to the username into the **Password** field.
- Step 4. Enter the name of the share into the **Share** field.
SSB saves all data into this directory, automatically creating the subdirectories. Backups of log files are stored in the *data*, configuration backups in the *config* subdirectory.
- Step 5. Enter the domain name of the target server into the **Domain** field.
- Step 6. Click **Commit**.

4.7.4.3. Procedure – Configuring NFS

Purpose:

The **NFS** backup method connects to a shared directory of the target server with the Network File Share protocol. To configure this method, enter the name of the NFS export into the **Export** field. SSB saves all data into this directory, automatically creating the subdirectories.



Figure 4.25. Configuring NFS backups

The backup server must also be configured to accept backups from SSB. To configure NFS on the remote server, complete the following steps:

**Note**

These steps must be performed on the remote server, not on SSB.

Steps:

- Step 1. Add a line that corresponds to the settings of SSB to the `/etc/exports` file of the backup server. This line should contain the following parameters:
- The path to the backup directory as set in the **Export** field of the SSB backup or archiving policy.
 - The IP address of the SSB interface that is used to access the remote server (that is, the address of the external interface, or the address of the management interface if it is enabled and the routing table of SSB is correctly configured — for details, see *Section 4.3, Network settings* (p. 37).
 - The following parameters: `(rw, no_root_squash, sync)`.

**Example 4.4. Configuring NFS on the remote server**

For example, if SSB connects the remote server from the `192.168.1.15` IP address and the data is saved into the `/var/backups/SSB` directory, add the following line to the `/etc/exports` file:

```
/var/backups/SSB 192.168.1.15(rw,no_root_squash, sync)
```

Step 2. Execute the following command: `exportfs -a`

Step 3. Verify that the `rpc portmapper` and `rpc.statd` applications are running.

4.7.5. Ownership of the backup files

The different backup protocols assign different file ownerships to the files saved on the backup server. The owners of the backup files created using the different protocols are the following:

- **rsync:** The user provided on the web interface.
- **SMB:** The user provided on the web interface.
- **NFS:** `root` with `no-root-squash`, `nobody` otherwise.

**Warning**

SSB cannot modify the ownership of a file that already exists on the remote server. If you change the backup protocol but you use the same directory of the remote server to store the backups, make sure to adjust the ownership of the existing files according to the new protocol. Otherwise SSB cannot overwrite the files and the backup procedure fails.



Chapter 5. User management and access control

The **AAA** menu (Authentication, Authorization, and Accounting) allows you to control the authentication, authorization, and accounting settings of the users accessing SSB. The following will be discussed in the next sections:

- For details on how to authenticate locally on SSB — see *Procedure 5.1, Managing SSB users locally* (p. 69).
- For details on how to authenticate users using an external LDAP (for example Microsoft Active Directory) database — see *Procedure 5.4, Managing SSB users from an LDAP database* (p. 73).
- For details on how to authenticate users using an external RADIUS server — see *Procedure 5.5, Authenticating users to a RADIUS server* (p. 76).
- For details on how to control the privileges of users and usergroups — see *Section 5.6, Managing user rights and usergroups* (p. 77).
- For details on how to display the history of changes of SSB configuration — see *Section 5.7, Listing and searching configuration changes* (p. 83).

5.1. Procedure – Managing SSB users locally

Purpose:

By default, SSB users are managed locally on SSB. To create and delete local users, modify the group membership of local users, or to modify the password of a user, complete the following procedure.



Note

The *admin* user is available by default and has all possible privileges. It is not possible to delete this user.

Local users cannot be managed when LDAP authentication is used (see *Procedure 5.4, Managing SSB users from an LDAP database* (p. 73)). When LDAP authentication is enabled, the accounts of local users is disabled, they are not displayed on the **AAA > Local Users** page, but they are not deleted,

When using RADIUS authentication together with local users, the users are authenticated to the RADIUS server, only their group memberships must be managed locally on SSB. For details, see *Procedure 5.5, Authenticating users to a RADIUS server* (p. 76).

Steps:

Step 1. Navigate to **AAA > Local Users** and click **+**.



Figure 5.1. Creating local users

Step 2. Enter the username into the **User** field.



Note

The following characters cannot be used in usernames: \/[] ; ; | = , + * ? < >

Step 3. Enter a password for the user into the **Password** and **Verify password** fields.

The strength of the password is indicated below the **Password** field as you type. To set a policy for password strength, see *Procedure 5.2, Setting password policies for local users (p. 70)*. The user can change the password later from the SSB web interface.

Step 4. Click **+** in the **Groups** section and select a group that the user will be member of. Repeat this step to add the user to multiple groups. For details about the different groups, see *Section 5.6, Managing user rights and usergroups (p. 77)*.

- To remove a user from a group, click **x** next to the group.
- To delete a user, click **x** at the right edge of the screen.

Step 5. Click **Commit**.

5.2. Procedure – Setting password policies for local users

Purpose:

SSB can use password policies to enforce minimal password strength and password expiry. To create a password policy, complete the following steps.



Note

Password policies apply only for locally managed users, it has no effect if you manage your users from an LDAP database, or if you authenticate your users to a RADIUS server.

Password policies do not apply to the built-in *admin* user.

Steps:

Step 1. Navigate to **AAA > Settings**.



The screenshot shows the SSB configuration interface. On the left, there is a 'Main menu' with options like Basic Settings, AAA, Policies, Log, Search, and Reports. Below it is a 'User menu' with Private keystore, Change password, Preferences, and Logout. Further down is 'User info' showing details for user 'admin'. At the bottom left is a 'System monitor' showing time and HA status. The main area is titled 'Settings' and has tabs for Group Management, Local Users, Access Control, and Accounting. The 'Local Users' tab is active, showing 'Authentication settings'. Under 'Authentication method', 'Password provided by database' is selected. Under 'User database', 'Local' is selected. Password policy fields include 'Password expiration' (0 days), 'Number of passwords to remember' (10), and 'Minimal password strength' (good). There is also a 'Cracklib' checkbox which is disabled. At the bottom, there is an 'Accounting settings' section with a 'Require commit log' checkbox. 'Commit' buttons are present in the top right and bottom right corners.

Figure 5.2. Configuring password policies

Step 2. Verify that the **Authentication method** is set to **Password provided by database** and that the **User database** is set to **Local**.



Note

If the setting of these fields is different (for example LDAP or RADIUS), then SSB is not configured to manage passwords locally.

Step 3. Set how long the passwords are valid in the **Password expiration** field. After this period, SSB users will have to change their password. To disable password expiry, enter 0.

Step 4. To prevent password-reuse (for example when a user has two password and instead of changing to a new password only switches between the two), set how many different passwords must the user use before reusing an old password.

Step 5. To enforce the use of strong password, select the level of password-complexity from the **Minimal password strength** field.



Note

The strength of the password is determined by its entropy: the variety of numbers, letters, capital letters, and special characters used, not only by its length.

The **Enable cracklib** option executes some simple dictionary-based attacks to find weak passwords.

Step 6. Click **Commit**.



Note

Changes to the password policy do not affect existing passwords. However, setting password expiry will require every user to change their passwords after the expiry date, and the new passwords must comply with the strength requirements set in the password policy.

5.3. Procedure – Managing local usergroups

Purpose:

To display which users belong to a particular local usergroup, navigate to **AAA > Group Management**. You can edit the group memberships here as well.

You can use local groups to control the privileges of SSB local and LDAP users — who can view and configure what. Local groups can be also used to control access to the logfiles available via a shared folder. For details, see *Section 8.6, Accessing log files across the network (p. 128)*.

For the description of built-in groups, see *Section 5.6.5, Built-in usergroups of SSB (p. 81)*. To create a new group, complete the following steps:

Steps:

Step 1. Navigate to **AAA > Group Management** and click **+**.

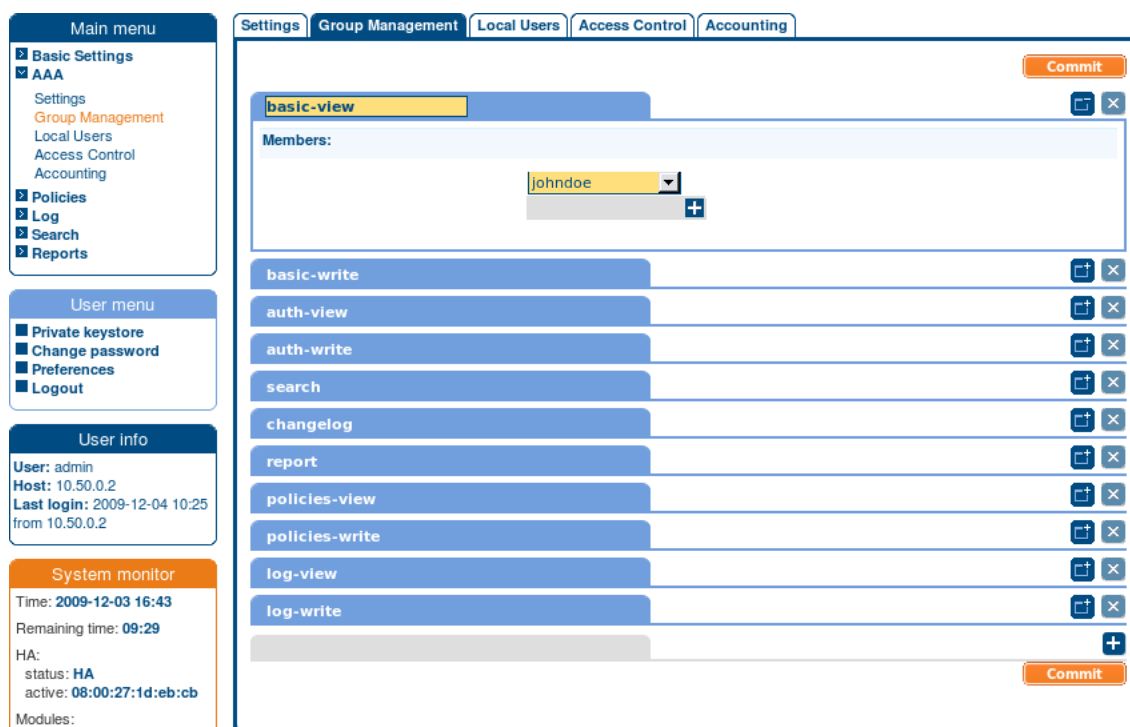


Figure 5.3. Group management

Step 2. Enter a name for the group.

Step 3. Enter the names of the users belonging to the group. Click **+** to add more users.



Step 4. Click **Commit**.

5.4. Procedure – Managing SSB users from an LDAP database

Purpose:

The SSB web interface can authenticate users to an external LDAP database to simplify the integration of SSB to your existing infrastructure. You can also specify multiple LDAP servers; if the first server is unavailable, SSB will try to connect to the second server. To enable LDAP authentication, complete the following steps.



Note

The *admin* user is available by default and has all privileges. It is not possible to delete this user.

The *admin* user can login to SSB even if LDAP authentication is used.

Enabling LDAP authentication automatically disables the access of every local user except for *admin*.

SSB accepts both pre-win2000-style and Win2003-style account names (User Principal Names). User Principal Names (UPNs) consist of a username, the at (@) character, and a domain name, for example *administrator@example.com*.

The following characters cannot be used in usernames and group names: `/\[]:;|=, + *) ? < > @`

When using RADIUS authentication together with LDAP users, the users are authenticated to the RADIUS server, only their group memberships must be managed in LDAP. For details, see *Procedure 5.5, Authenticating users to a RADIUS server* (p. 76).



Warning

A user can belong to a maximum of 10,000 groups; further groups are ignored.

Steps:


Step 1. Navigate to **AAA > Settings > Authentication settings**.

Step 2. Select the **LDAP** option and enter the parameters of your LDAP server.



Figure 5.4. Configuring LDAP authentication

Step a. Enter the IP address or hostname and port of the LDAP server into the **Server Address** field. If you want to encrypt the communication between SSB and the LDAP server, in case of SSL/TLS, enter 636 as the port number, or in case of STARTTLS, enter 389 as the port number.

To add multiple servers, click  and enter the address of the next server. If a server is unreachable, SSB will try to connect to the next server in the list in failover fashion.



Warning

If you will use a TLS-encrypted with certificate verification to connect to the LDAP server, use the full domain name (for example *ldap.example.com*) in the **Server Address** field, otherwise the certificate verification might fail. The name of the LDAP server must appear in the Common Name of the certificate.

Step b. Select the type of your LDAP server in the **Type** field. Select **Active Directory** to connect to Microsoft Active Directory servers, or **Posix** to connect to servers that use the POSIX LDAP scheme.

Step c. Enter the name of the DN to be used as the base of the queries into the **Base DN** field (for example *DC=demodomain*, *DC=exampleinc*).

Step d. Enter the name of the DN where SSB should bind to before accessing the database into the **Bind DN** field.

For example: *CN=Administrator*, *CN=Users*, *DC=demodomain*, *DC=exampleinc*.

**Note**

SSB accepts both pre-win2000-style and Win2003-style account names (User Principal Names), for example *administrator@example.com* is also accepted.

Step e. Enter the password to use when binding to the LDAP server into the **Bind Password** field.

Step 3. If you want to encrypt the communication between SSB and the LDAP server, in **Encryption**, select the **SSL/TLS** or the **STARTTLS** option and complete the following steps:

**Note**

TLS-encrypted connection to Microsoft Active Directory is supported only on Windows 2003 Server and newer platforms. Windows 2000 Server is not supported.

- If you want SSB to verify the certificate of the server, select **Only accept certificates authenticated by the specified CA certificate** and click the ☒ icon in the **CA X.509 certificate** field. A popup window is displayed.

Click **Browse**, select the certificate of the Certificate Authority (CA) that issued the certificate of the LDAP server, then click **Upload**. Alternatively, you can paste the certificate into the **Copy-paste** field and click **Set**.

SSB will use this CA certificate to verify the certificate of the server, and reject the connections if the verification fails.

**Warning**

If you will use a TLS-encrypted with certificate verification to connect to the LDAP server, use the full domain name (for example *ldap.example.com*) in the **Server Address** field, otherwise the certificate verification might fail. The name of the LDAP server must appear in the Common Name of the certificate.

- If the LDAP server requires mutual authentication, that is, it expects a certificate from SSB, enable **Authenticate as client**. Generate and sign a certificate for SSB, then click ☒ in the **Client X.509 certificate** field to upload the certificate. After that, click ☒ in the **Client key** field and upload the private key corresponding to the certificate.

Step 4. Click **Commit**.

**Note**

You also have to configure the usergroups in SSB and possibly in your LDAP database. For details on using usergroups, see *Section 5.6.4, How to use usergroups (p. 80)*.



Step 5. Click **Test** to test the connection. Note that the testing of SSL-encrypted connections is currently not supported.

5.5. Procedure – Authenticating users to a RADIUS server

Purpose:

SSB can authenticate its users to an external RADIUS server. Group memberships of the users must be managed either locally on SSB or in an LDAP database.



Warning

The challenge/response authentication methods is currently not supported. Other authentication methods (for example password, SecureID) should work.

To authenticate SSB users to a RADIUS server, complete the following steps:

Steps:

Step 1. Navigate to **AAA > Settings**.

The screenshot displays the SSB configuration interface. On the left, there is a sidebar with a 'Main menu' containing 'Basic Settings', 'AAA', 'Settings', 'Group Management', 'Local Users', 'Access Control', and 'Accounting'. Below this is a 'User menu' with 'Private keystore', 'Change password', 'Preferences', and 'Logout'. At the bottom of the sidebar is 'User info' showing 'User: admin', 'Host: 10.50.0.2', and 'Last login: 2009-12-04 10:56 from 10.50.0.2'. Below that is a 'System monitor' section showing 'Time: 2009-12-03 16:43', 'Remaining time: 09:29', 'HA: status: HA, active: 08:00:27:1d:eb:cb', 'Modules: syslog-ng: RUNNING', 'Active: Hosts: UNKNOWN, Senders: UNKNOWN', and 'Load 1: 0.00 Load 15: 0.02'. The main content area is titled 'Settings' and has tabs for 'Group Management', 'Local Users', 'Access Control', and 'Accounting'. The 'Access Control' tab is selected, showing 'Authentication settings'. Under 'Authentication method:', 'RADIUS' is selected. The 'Server address:' section has a table with columns 'Address', 'Port', and 'Shared secret'. The first row contains 'radius.example.com', '1812', and '*****'. There is a '+' button to add more servers. Below this, 'User database:' has 'Local' selected. The 'Password expiration:' is set to '0' days, 'Number of passwords to remember:' is '10', and 'Minimal password strength:' is 'good'. There is a 'Cracklib (eg. dictionary) check on password:' section with 'Enabled' selected. At the bottom, 'Accounting settings' has a 'Require commit log:' checkbox which is unchecked. There are 'Commit' buttons in the top right and bottom right of the main content area.

Figure 5.5. Configuring RADIUS authentication

Step 2. Set the **Authentication method** field to **RADIUS**.

Step 3. Enter the IP address or domain name of the RADIUS server into the **Address** field.

Step 4. Enter the password that SSB can use to access the server into the **Shared secret** field.

Step 5. To add more RADIUS servers, click **+** and repeat Steps 2-4.



Repeat this step to add multiple servers. If a server is unreachable, SSB will try to connect to the next server in the list in failover fashion.

Step 6. When configuring RADIUS authentication with a local user database, complete the following steps.

Step a. Set **Password expiration** to 0.

Step b. Set **Number of passwords to remember** to 0.

Step c. Set **Minimal password strength** to *disabled*.

Step d. Set **Cracklib check on password** to *disabled*.

Step 7.



Warning

After clicking **Commit**, the SSB web interface will be available only after successfully authenticating to the RADIUS server. Note that the default *admin* account of SSB will be able to login normally, even if the RADIUS server is inaccessible.

Click **Commit**.

5.6. Managing user rights and usergroups

In SSB, user rights can be assigned to usergroups. SSB has numerous usergroups defined by default, but custom user groups can be defined as well. Every group has a set of privileges: which pages of the SSB web interface it can access, and whether it can only view (read) or also modify (read & write/perform) those pages or perform certain actions.



Note

Every group has either read or read & write/perform privileges to a set of pages.

- For details on modifying existing groups, see *Procedure 5.6.1, Modifying group privileges (p. 78)*.
- For details on creating a new usergroup, see *Procedure 5.6.2, Creating new usergroups for the SSB web interface (p. 79)*.
- For details on finding usergroups that have a specific privilege, see *Section 5.6.3, Finding specific usergroups (p. 80)*.
- For tips on using usergroups, see *Section 5.6.4, How to use usergroups (p. 80)*.
- For a detailed description about the privileges of the built-in usergroups, see *Section 5.6.5, Built-in usergroups of SSB (p. 81)*.

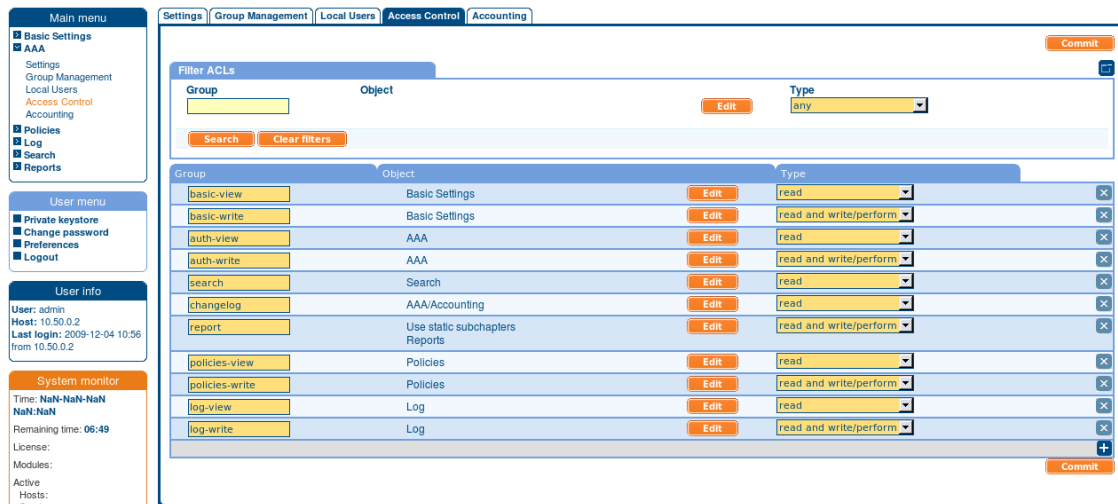


Figure 5.6. Managing SSB users

5.6.1. Procedure – Modifying group privileges

Purpose:

To modify the privileges of an existing group, complete the following steps:

Steps:

- Step 1. Navigate to **AAA > Access Control**.
- Step 2. Find the group you want to modify and click . The list of available privileges is displayed.
- Step 3. Select the privileges (pages of the SSB interface) to which the group will have access to and click **Save**.

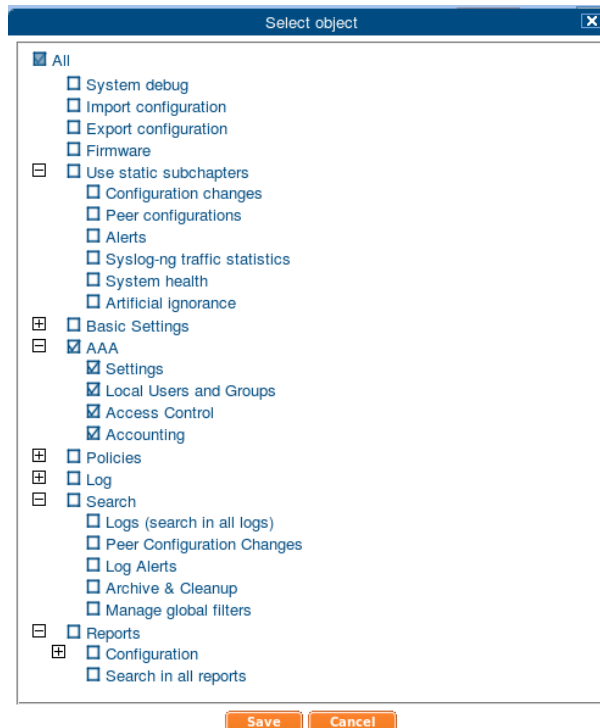


Figure 5.7. Modifying group privileges



Warning

Assigning the **Search** privilege to a user on the AAA page grants the user search access to every logspace, even if the user is not a member of the groups listed in the **Access control** option of the particular logspace.

Step 4. Select the type of access (read or read & write) from the **Type** field.


Step 5. Click **Commit**.

5.6.2. Procedure – Creating new usergroups for the SSB web interface

Purpose:

To create a new group, complete the following steps:

Steps:

- Step 1. Navigate to **AAA > Access Control** and click **+**.
- Step 2. Enter a name for the group. For details on how you should name your groups, see *Section 5.6.4, How to use usergroups (p. 80)*.
- Step 3. Click the  icon located next to the name of the group. The list of available privileges is displayed.
- Step 4. Select the privileges (pages of the SSB interface) to which the group will have access to and click **Save**.

**Note**

To export the configuration of SSB, the **Export configuration** privilege is required.

To import a configuration to SSB, the **Import configuration** privilege is required.

To update the firmware and set the active firmware, the **Firmware** privilege is required.

Step 5. Select the type of access (read or read & write) from the **Type** field.

Step 6. Click **Commit**.

The *admin* user is available by default and has all privileges, except that it cannot remotely access the shared logspaces. It is not possible to delete this user.

5.6.3. Finding specific usergroups

The **Filter ACLs** section of the **AAA > Access Control** page provides you with a simple searching and filtering interface to search the names and privileges of usergroups.

Group	Object	Type
basic-view	Basic Settings	read
basic-write	Basic Settings	read and write/performance
auth-view	AAA	read
auth-write	AAA	read and write/performance
search	Search	read
changelog	AAA/Accounting	read
report	Use static subchapters Reports	read and write/performance
policies-view	Policies	read
policies-write	Policies	read and write/performance
log-view	Log	read
log-write	Log	read and write/performance

Figure 5.8. Finding specific usergroups

- To select usergroups starting with a specific string, enter the beginning of the name of the group into the **Group** field and select **Search**.
- To select usergroups who have a specific privilege, click ☒, select the privilege or privileges you are looking for, and click **Search**.
- To filter for read or write access, use the **Type** option.

5.6.4. How to use usergroups

How you should name usergroups depends on the way you manage your SSB users.

- **Local users:** If you use only local users, create or modify your usergroups on the **AAA > Access Control** page and add users to the groups on the **AAA > Local Users** or the **AAA > Group Management** page.



- **LDAP users and LDAP groups:** If you manage your users from LDAP, and also have LDAP groups that match the way you want to group your SSB users, create or modify your usergroups on the **AAA > Access Control** page and ensure that the name of your LDAP group and the SSB usergroup is the same. For example, to make members of the *admins* LDAP group be able to use SSB, create a usergroup called *admins* on the **AAA > Access Control** page and edit the privileges of the group as needed.



Warning

A user can belong to a maximum of 10,000 groups; further groups are ignored.

- **RADIUS users and local groups:** This is the case when you manage users from RADIUS, but you cannot or do not want to create groups in LDAP. Create your local groups on the **AAA > Access Control** page, and add your RADIUS users to these groups on the **AAA > Group Management** page.

5.6.5. Built-in usergroups of SSB

SSB has the following usergroups by default:

Group	Object	Type
basic-view	Basic Settings	read
basic-write	Basic Settings	read and write/perform
auth-view	AAA	read
auth-write	AAA	read and write/perform
search	Search	read
changelog	AAA/Accounting	read
report	Use static subchapters Reports	read and write/perform
policies-view	Policies	read
policies-write	Policies	read and write/perform
log-view	Log	read
log-write	Log	read and write/perform

Figure 5.9. Built-in usergroups of SSB



Warning

If you use LDAP authentication on the SSB web interface and want to use the default usergroups, you have to create these groups in your LDAP database and assign users to them. For details on using usergroups, see *Section 5.6.4, How to use usergroups (p. 80)*.

- **basic-view:** View the settings in the **Basic Settings** menu, including the system logs of SSB. Members of this group can also execute commands on the **Troubleshooting** tab.



- **basic-write**: Edit the settings in the **Basic Settings** menu. Members of this group can manage SSB as a host.
- **auth-view**: View the names and privileges of the SSB administrators, the configured usergroups, and the authentication settings in the **AAA** menu. Members of this group can also view the history of configuration changes.
- **auth-write**: Edit authentication settings and manage users and usergroups.

**Warning**

Members of the *auth-write* group, or any other group with write privileges to the **AAA** menu are essentially equivalent to system administrators of SSB, because they can give themselves any privilege. Users with limited rights should never have such privileges.

If a user with write privileges to the **AAA** menu gives himself new privileges (for example gives himself group membership to a new group), then he has to relogin to the SSB web interface to activate the new privilege.

- **search**: Browse and download various logs and alerts in the **Search** menu.

**Note**

The *admin* user is not a member of this group by default, so it cannot remotely access the shared logspaces.

- **changelog**: View the history of SSB configuration changes in the **AAA > Accounting** menu.
- **report**: Browse, create and manage reports, and add statistics-based chapters to the reports in the **Reports** menu.

**Note**

To control exactly which statistics-based chapters and reports can the user include in a report, use the *Use static subchapters* privileges.

- **policies-view**: View the policies and settings in the **Policies** menu.
- **policies-write**: Edit the policies and settings in the **Policies** menu.

**Warning**

Members of this group can make the logs stored on SSB available as a shared network drive. In case of unencrypted logfiles, this may result in access to sensitive data.

- **log-view**: View the logging settings in the **Log** menu.
- **log-write**: Configure logging settings in the **Log** menu.



5.7. Listing and searching configuration changes

SSB automatically tracks every change of its configuration. To display the history of changes, select **AAA > Accounting**. The changes are organized as log messages, and can be browsed and searched using the regular SSB search interface (for details, see *Chapter 12, Browsing log messages and SSB reports (p. 156)*). The following information is displayed about each modification:

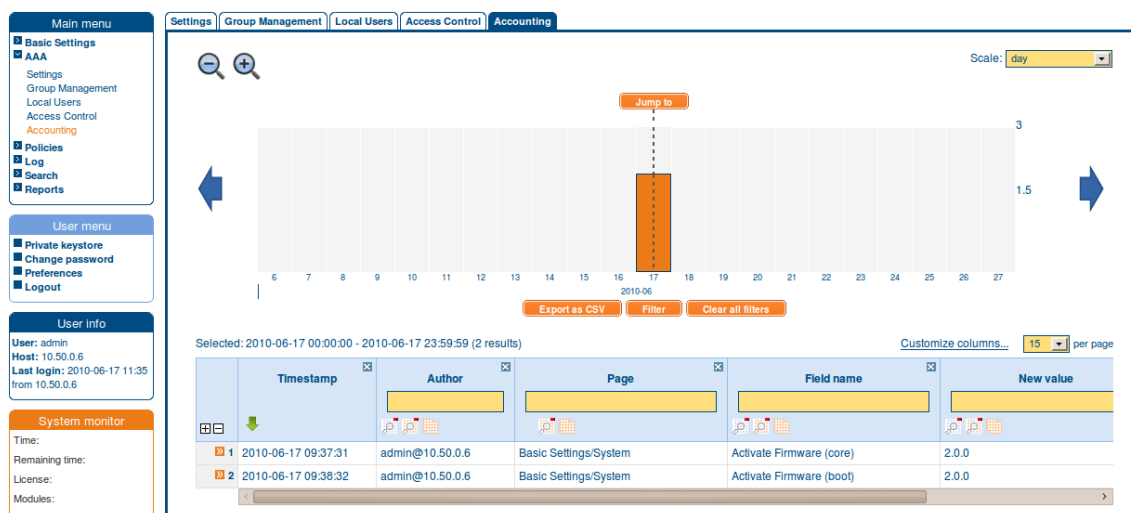


Figure 5.10. Browsing configuration changes

- **Timestamp:** The date of the modification.
- **Author:** Username of the administrator who modified the configuration of SSB.
- **Page:** The menu item that was modified.
- **Field name:** The name of the field or option that was modified.
- **New value:** The new value of the configuration parameter.
- **Message:** The changelog or commit log that the administrator submitted. This field is available only if the **Require commit log** option is enabled (see below).
- **Old value:** The old value of the configuration parameter.

To request the administrators to write an explanation to every configuration change, navigate to **AAA > Settings > Accounting settings** and select the **Require commit log** option.



Chapter 6. Managing SSB

The following sections explain the basic management tasks of SSB, including the basic control (for example, shutdown or reboot) of the appliance, upgrading, as well as tips on troubleshooting SSB.

6.1. Controlling SSB — restart, shutdown

To restart or shut down SSB, navigate to **Basic Settings** > **System** > **System control** > **This node** and click the respective action button. The **Other node** refers to the slave node of a high availability SSB cluster. For details on high availability clusters, see *Section 6.2, Managing a high availability SSB cluster (p. 84)*.



Warning

- When rebooting the nodes of a cluster, reboot the other (slave) node first to avoid unnecessary takeovers.
- When shutting down the nodes of a cluster, shut down the other (slave) node first. When powering on the nodes, start the master node first to avoid unnecessary takeovers.
- When both nodes are running, avoid interrupting the connection between the nodes: do not unplug the Ethernet cables, reboot the switch or router between the nodes (if any), or disable the HA interface of SSB.




Figure 6.1. Performing basic management



Note

Web sessions to the SSB interface are persistent and remain open after rebooting SSB, so you do not have to relogin after a reboot.

6.2. Managing a high availability SSB cluster



Warning

- When rebooting the nodes of a cluster, reboot the other (slave) node first to avoid unnecessary takeovers.
- When shutting down the nodes of a cluster, shut down the other (slave) node first. When powering on the nodes, start the master node first to avoid unnecessary takeovers.
- When both nodes are running, avoid interrupting the connection between the nodes: do not unplug the Ethernet cables, reboot the switch or router between the nodes (if any), or disable the HA interface of SSB.

The **Basic Settings** > **High Availability** page provides the following information about SSB:

**Note**

Refer to *Procedure B.2, Installing two SSB units in HA mode* (p. 205) of *Appendix B, syslog-ng Store Box Hardware Installation Guide* (p. 204) for details on creating a high availability SSB cluster.

The screenshot shows the SSB web interface with the following components:

- Main menu:** Basic Settings, Network, System, High Availability, Date & Time, Management, Alerting & Monitoring, Troubleshooting, Dashboard, AAA, Policies, Log, Search, Reports.
- User menu:** Private keystore, Change password, Preferences, Logout.
- User info:** User: admin, Host: 10.50.0.2, Last login: 2009-12-04 10:56 from 10.50.0.2.
- System monitor:** Time: 2009-12-03 16:43, Remaining time: 09:29, HA status: HA, active: 08:00:27:1d:eb:cb, Modules: syslog-ng: RUNNING, Active Hosts: UNKNOWN, Senders: UNKNOWN, Load 1: 0.00, Load 15: 0.02, CPU, Mem, Disk, Swap usage bars.
- High Availability & Nodes:**
 - Status: SSB is currently operating in HA state.
 - Current master: 08:00:27:1d:eb:cb (Buttons: Activate Slave, Synchronize configuration, Reboot cluster)
 - HA UUID: bc7870f9-d022-409e-bdfe-29c4e51c59b7
 - DRBD Status: Connected - Connected, Connected
 - DRBD sync rate limit: 10 Mba/s
 - This node:**
 - Node ID: 08:00:27:1d:eb:cb
 - Node HA state: HA
 - Node HA UUID: bc7870f9-d022-409e-bdfe-29c4e51c59b7
 - DRBD status: Connected (UpToDate)
 - RAID status: Data partition: active
 - Boot firmware versions: Current: HEAD1259758326, Active: HEAD1259758326
 - HA link speed: Auto negotiation
 - Interfaces for Heartbeat: HA (Fix current), External, Internal, Management
 - Next hop monitoring: External, Internal, Management
 - Buttons: Reboot, Shutdown
 - Other node:**
 - Node ID: 08:00:27:a1:2f:52
 - Node HA state: HA
 - Node HA UUID: bc7870f9-d022-409e-bdfe-29c4e51c59b7
 - DRBD status: Connected (UpToDate)
 - RAID status: Data partition: active
 - Boot firmware versions: Current: HEAD1259758326, Active: HEAD1259758326
 - HA link speed: Auto negotiation
 - Interfaces for Heartbeat: HA (Fix current), External, Internal, Management
 - Next hop monitoring: External, Internal, Management
 - Buttons: Reboot, Shutdown

Figure 6.2. Managing a high availability cluster

- **Status:** Indicates whether SSB is running in High Availability or Standalone mode.
- **Node ID:** The MAC address of the HA interface of the node. This address is also printed on a label on the top cover of the SSB unit.
- **Node HA state:** Indicates whether the SSB node is running in High Availability or Standalone mode.
- **Node HA UUID:** A unique identifier of the node. Only available in High Availability mode.
- **DRBD status:** The status of data stored on SSB. The status must be *Consistent* on the active node to prevent data loss.
- **RAID status:** The status of the RAID device of the node.

The active (master) SSB node is labeled as **This node**, this unit receives the incoming log messages and provides the web interface. The SSB unit labeled as **Other node** is the slave node that is activated if the master node becomes unavailable.

**Note**

For SSB clusters, the ID of the node (the MAC address of its HA interface) sending the message is included in the log messages.



To activate the other node and disable the currently active node, click **Activate slave**.



Warning

Hazard of data loss! Activating the slave node terminates all connections of SSB and might result in data loss. The slave node becomes active after about 60 seconds, during which SSB cannot accept incoming messages. Enable disk-buffering on your syslog-ng clients and relays to prevent data loss in such cases.

To reboot both nodes, click **Reboot Cluster**. To prevent takeover, a token is placed on the slave node. While this token persists, the slave node halts its boot process to make sure that the master node boots first. Following reboot, the master removes this token from the slave node, allowing it to continue with the boot process.

If the token still persists on the slave node following reboot, the **Unblock Slave Node** button is displayed. Clicking the button removes the token, and reboots the slave node.

6.2.1. High Availability status explained

This section explains the possible statuses of the SSB nodes and the DRBD data storage system. SSB displays this information on the **Basic Settings > High Availability** page.



Note

If a redundant Heartbeat interface is configured, its status is also displayed in the **Redundant Heartbeat status** field. For details on redundant Heartbeat interfaces, see *Section 6.2.4, Redundant Heartbeat interface status explained (p. 90)* and *Procedure 6.2.3, Configuring redundant Heartbeat interfaces (p. 89)*.

The **Status** field indicates whether the SSB nodes recognize each other properly and whether those are configured to operate in high availability mode. The status of the individual SSB nodes is indicated in the **Node HA status** field of the each node. The following statuses can occur:

- **Standalone:** There is only one SSB unit running in *standalone* mode, or the units have not been converted to a cluster (the **Node HA status** of both nodes is *standalone*). Click **Convert to Cluster** to enable High Availability mode.
- **HA:** The two SSB nodes are running in High Availability mode. **Node HA status** is *HA* on both nodes, and the **Node HA UUID** is the same on both nodes.
- **Half:** High Availability mode is not configured properly, one node is in *standalone*, the other one in *HA* mode. Connect to the node in *HA* mode, and click **Join HA** to enable High Availability mode.
- **Broken:** The two SSB nodes are running in High Availability mode. **Node HA status** is *HA* on both nodes, but the **Node HA UUID** is different. Contact the BalaBit Support Team for help. For contact details, see *Section 5, Contact and support information (p. xiii)*.
- **Degraded:** SSB was running in high availability mode, but one of the nodes has disappeared (for example broken down, or removed from the network). Power on, reconnect, or repair the missing node.
- **Degraded Sync:** Two SSB units were joined to High Availability mode, and the first-time synchronization of the disks is currently in progress. Wait for the synchronization to complete. Note that in case of large disks with lots of stored data, synchronizing the disks can take several hours.



- **Split brain:** The two nodes lost the connection to each other, with the possibility of both nodes being active (master) for a time.



Warning

Hazard of data loss! In this case, valuable log messages might be available on both SSB nodes, so special care must be taken to avoid data loss. For details on solving this problem, see *Procedure 14.6.2, Recovering from a split brain situation (p. 198)*.

Do NOT reboot or shut down the nodes.

- **Invalidated:** The data on one of the nodes is considered out-of-sync and should be updated with data from the other node. This state usually occurs during the recovery of a split-brain situation when the DRBD is manually invalidated.
- **Converted:** After converting nodes to a cluster (clicking **Convert to Cluster**) or enabling High Availability mode (clicking **Join HA**) and before rebooting the node(s).



Note

If you experience problems because the nodes of the HA cluster do not find each other during system startup, navigate to **Basic Settings > High Availability** and select **Make HA IP permanent**. That way the IP address of the HA interfaces of the nodes will be fixed, which helps if the HA connection between the nodes is slow.

The **DRBD status** field indicates whether the latest data (including SSB configuration, log files, and so on) is available on both SSB nodes. The master node (this node) must always be in **consistent** status to prevent data loss. Inconsistent status means that the data on the node is not up-to-date, and should be synchronized from the node having the latest data.

The **DRBD status** field also indicates the connection between the disk system of the SSB nodes. The following statuses are possible:

- **Connected:** Both nodes are functioning properly.
- **Invalidated:** The data on one of the nodes is considered out-of-sync and should be updated with data from the other node. This state usually occurs during the recovery of a split-brain situation when the DRBD is manually invalidated.
- **Sync source or Sync target:** One node (**Sync target**) is downloading data from the other node (**Sync source**).



Note

When the two nodes are synchronizing data, it is not possible to reboot or shutdown the master node. If you absolutely must shutdown SSB in such a situation, shutdown the slave node first, and then the master node.

When synchronizing data, the progress and the remaining time is displayed in the **System monitor**.



- **Split brain:** The two nodes lost the connection to each other, with the possibility of both nodes being active (master) for a time.



Warning

Hazard of data loss! In this case, valuable log messages might be available on both SSB nodes, so special care must be taken to avoid data loss. For details on solving this problem, see *Procedure 14.6.2, Recovering from a split brain situation* (p. 198).

- **WFConnection:** One node is waiting for the other node; the connection between the nodes has not been established yet.

6.2.2. Adjusting the synchronization speed of DRBD

When operating two SSB units in High Availability mode, every incoming data copied from the master (active) node to the slave (passive) node. Since synchronizing data can take up significant system-resources, the maximal speed of the synchronization is limited, by default, to **10 MB/sec**. However, this means that synchronizing large amount of data can take very long time, so it is useful to increase the synchronization speed in certain situations — for example, when synchronizing the disks after converting a single node to a high availability cluster.

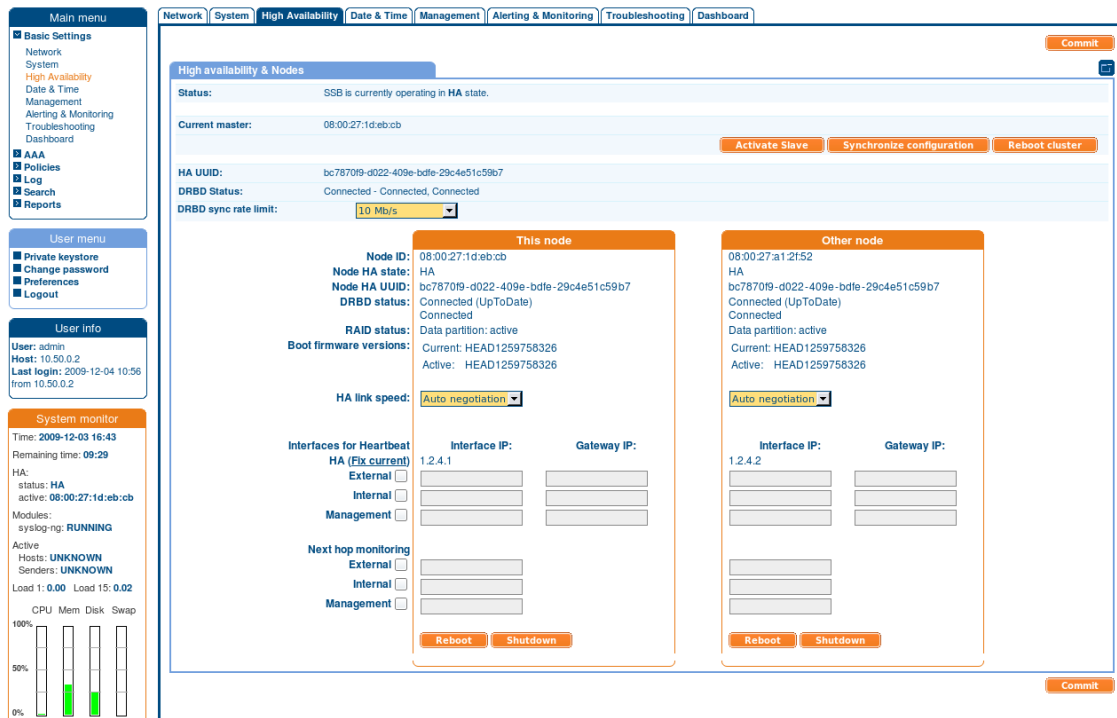


Figure 6.3. Adjusting DRBD synchronization speed

To change the limit of the DRBD synchronization rate, navigate to **Basic Settings > High Availability**, select **DRBD sync rate limit**, and select the desired value.

**Note**

Setting the sync rate to a high value is not recommended if the load of SSB is very high, because increasing the resources used by the synchronization process may degrade the general performance of SSB.

6.2.3. Procedure – Configuring redundant Heartbeat interfaces

Purpose:

In order to avoid unnecessary takeovers and minimize the chance of split-brain situations, it is possible to configure additional Heartbeat interfaces in SSB. These redundant Heartbeat interfaces are used only to detect that the other node is still available; it is not used to synchronize data between the nodes (only Heartbeat messages are transferred). For example, if the main HA interface breaks down, or is accidentally unplugged and the nodes can still access each other on the redundant HA interface, no takeover occurs, but no data is synchronized to the slave node until the main HA link is restored. Similarly, if connection on the redundant Heartbeat interface is lost, but the main HA connection is available, no takeover occurs.

The redundant Heartbeat interface is a virtual interface that uses an existing interface of the SSB device (for example the external or the management interface). The original MAC address of the interface is displayed at **Basic Settings > High Availability > Interfaces for Heartbeat > Production MAC**, while the MAC address of the virtual redundant Heartbeat interface is displayed at **Basic Settings > High Availability > Interfaces for Heartbeat > HA MAC**. The MAC address of the redundant Heartbeat interface is generated in a way that it cannot interfere with the MAC addresses of physical interfaces.

**Warning**

In case the nodes lose connection on the main HA interface, and after a time the connection is lost on the redundant Heartbeat interfaces as well, the slave node will become active. However, as the master node was active for a time when no data synchronization was possible between the nodes, this results in a split-brain situation which must be resolved before the HA functionality can be restored. For details, see *Procedure 14.6.2, Recovering from a split brain situation (p. 198)*.

**Note**

Even if redundant HA links are configured, if the dedicated HA link fails, the slave node will not be visible on the High Availability page anymore.

To configure a redundant Heartbeat interface, complete the following steps.

Steps:

- Step 1. Navigate to **Basic Settings > High Availability > Interfaces for Heartbeat**.
- Step 2. Select the interface you want to use as redundant Heartbeat interface (for example *External1*). Using an interface as a redundant Heartbeat interface does not affect the original traffic of the interface.



Main menu

- Basic Settings
 - Network
 - System
 - High Availability
 - Date & Time
 - Management
 - Alerting & Monitoring
 - Troubleshooting
 - Dashboard
- AAA
 - Policies
 - Log
 - Search
 - Reports
- User menu
 - Private keystore
 - Change password
 - Preferences
 - Logout
- User info

User: admin
Host: 10.50.0.2
Last login: 2009-12-04 10:56 from 10.50.0.2
- System monitor

Time: 2009-12-03 16:43
Remaining time: 09:29

HA:
status: HA
active: 08:00:27:1d:eb:cb

Modules:
syslog-ng: RUNNING

Active
Hosts: UNKNOWN
Senders: UNKNOWN

Load 1: 0.00 Load 15: 0.02

CPU Mem Disk Swap

100%
50%
0%

1% 34% 25% 0%

High availability & Nodes

Status: SSB is currently operating in HA state.

Current master: 08:00:27:1d:eb:cb

HA UUID: bc7870f9-d022-409e-bdfe-29c4e51c59b7

DRBD Status: Connected - Connected, Connected

DRBD sync rate limit: 10 Mb/s

Buttons: **Commit**, **Activate Slave**, **Synchronize configuration**, **Reboot cluster**

	This node	Other node
Node ID:	08:00:27:1d:eb:cb	08:00:27:a1:2f:52
Node HA state:	HA	HA
Node HA UUID:	bc7870f9-d022-409e-bdfe-29c4e51c59b7	bc7870f9-d022-409e-bdfe-29c4e51c59b7
DRBD status:	Connected (UpToDate)	Connected (UpToDate)
RAID status:	active	active
Boot firmware versions:	Current: HEAD1259758326 Active: HEAD1259758326	Current: HEAD1259758326 Active: HEAD1259758326
HA link speed:	Auto negotiation	Auto negotiation
Interfaces for Heartbeat	<p>Interface IP: 1.2.4.1 Gateway IP: 10.50.0.113</p> <p>Production MAC: 08:00:27:14:28:b5 HA MAC: 22:67:fd:1c:3e:6b</p> <p>Internal <input type="checkbox"/> Management <input type="checkbox"/></p> <p>Next hop monitoring External <input type="checkbox"/> Internal <input type="checkbox"/> Management <input type="checkbox"/></p> <p>Buttons: Reboot, Shutdown</p>	<p>Interface IP: 1.2.4.2 Gateway IP: 10.50.0.114</p> <p>Production MAC: 08:00:27:14:28:b5 HA MAC: fe:15:89:fa:36:a9</p> <p>Internal <input type="checkbox"/> Management <input type="checkbox"/></p> <p>Next hop monitoring External <input type="checkbox"/> Internal <input type="checkbox"/> Management <input type="checkbox"/></p> <p>Buttons: Reboot, Shutdown</p>

Commit

Figure 6.4. Configuring redundant Heartbeat interfaces

- Step 3. Enter an IP address into the **This node > Interface IP** field of the selected interface. This IP address must be a real IP address that is visible from the other node. The two nodes cannot have the same IP address on their redundant Heartbeat interfaces. Enter the IP address of the gateway into the **Gateway IP** if needed.
- Step 4. Enter an IP address into the **Other node > Interface IP** field of the selected interface. This IP address must be a real IP address that is visible from the other node. The two nodes cannot have the same IP address on their redundant Heartbeat interfaces. Enter the IP address of the gateway into the **Gateway IP** if needed.
- Step 5. Repeat the previous steps to add additional redundant Heartbeat interfaces if needed.
- Step 6. Click **Commit**.

**Warning**

For the changes to take effect, you have to restart both nodes. To restart both nodes, click **Reboot Cluster**.

6.2.4. Redundant Heartbeat interface status explained

The status of the redundant Heartbeat interfaces is displayed at **Basic Settings > High Availability > Redundant Heartbeat status**, and also in the **HA > Redundant** field of the System monitor. The possible status messages are explained below:



- **NOT USED:** There are no redundant Heartbeat interfaces configured.
- **OK:** Normal operation, every redundant Heartbeat interface is working properly.
- **DEGRADED-WORKING:** Two or more redundant Heartbeat interfaces are configured, and at least one of them is functioning properly. This status is displayed also when a new redundant Heartbeat interface has been configured, but the nodes of the SSB cluster has not been restarted yet.
- **DEGRADED:** The connection between the redundant Heartbeat interfaces has been lost. Investigate the problem to restore the connection.
- **INVALID:** An error occurred with the redundant Heartbeat interfaces. Contact the BalaBit Support Team for help. For contact details, see *Section 5, Contact and support information (p. xiii)*.

6.2.5. Procedure – Configuring next-hop router monitoring

Purpose:

By default, HA takeover occurs only if the master node stops working or becomes unreachable from the slave node. However, this does not cover the scenario when the master node becomes inaccessible to the outside world (for example its external interface or the router or switch connected to the external interface breaks down) while the slave node would be still accessible (for example because it is connected to a different router).

To address such situations, you can specify IP addresses (usually next hop routers) to continuously monitor from both the master and the slave nodes using ICMP echo (ping) messages. One such address can be set up for every interface.

When setting up next hop monitoring, you have to make sure that the master and slave nodes can ping the specified address directly. You can either:

- Choose the addresses of the redundant-HA SSB interfaces so that they are on the same subnet as the next-hop address
- Configure the next-hop device with an additional IP-address that is on the same subnet as the redundant-HA SSB interfaces facing it

If any of the monitored addresses becomes unreachable from the master node while being reachable from the slave node (in other words, more monitored addresses are accessible from the slave node) than it is assumed that the master node is unreachable and a forced takeover occurs — even if the master node is otherwise functional.

Naturally, if the slave node is not capable of taking over the master node (for example because there is data not yet synchronized from the current master node) no takeover is performed.

To configure a next hop monitoring, complete the following steps.

Steps:

- Step 1. Navigate to **Basic Settings > High Availability > Next hop monitoring**.
- Step 2. Select the interface to use for monitoring its next-hop router.



Figure 6.5. Configuring next hop monitoring

- Step 3. Enter the IP address to monitor from the current master node (for example the IP address of the router or the switch connected to the interface) into the **This node** > **Next hop IP** field of the selected interface. This IP address must be a real IP address that is visible from the interface, and must be on the same local network segment.
- Step 4. Enter the IP address to monitor from the current slave node (for example the IP address of the router or the switch connected to the interface) into the **Other node** > **Next hop IP** field of the selected interface. This IP address must be a real IP address that is visible from the interface, and must be on the same local network segment.
- Step 5. Repeat the previous steps to add IP addresses to be monitored from the other interfaces if needed.
- Step 6. Click **Commit**.

**Warning**

For the changes to take effect, you have to restart both nodes. To restart both nodes, click **Reboot Cluster**.

6.3. Upgrading SSB

The following sections describe how to keep SSB up to date, and how to install a new license file if needed.



- For details on how to upgrade SSB, see *Procedure 6.3.1, Updating SSB and managing the firmware* (p. 93).
- For details on how to upgrade the firmware of an SSB cluster, see *Procedure 6.3.2, Upgrading both the core and the boot firmware of a high availability system* (p. 95).
- For details on how to install a new license file, see *Procedure 6.3.4, Updating the SSB license* (p. 96).
- For details on how to import or export the configuration of SSB, see *Procedure 6.3.5, Exporting the configuration of SSB* (p. 97).

**Warning**

Downgrading from a feature release to an earlier (and thus unsupported) feature release, or to the stable release is not supported: this means that once you upgrade a system from a stable release (for example 1.0) to a feature release (for example 1.1), you will have to keep upgrading to the new feature releases until the next stable version release (for example 2.0) is published, or risk using an unsupported product.

6.3.1. Procedure – Updating SSB and managing the firmware

Purpose:

SSB can be updated when a new firmware version is available. To display information about the firmware currently running on SSB, navigate to **Basic Settings > System > Version details**. The following is displayed:

- **Core firmware | Boot firmware:** The version number of the firmwares currently running on SSB (for example 2.0).
- **Build date:** The date when the currently running firmware was created.

Updating SSB is described in the following sections.

Prerequisites:

**Warning**

Before uploading a new firmware image, backup the configuration of SSB. For details, see *Procedure 6.3.5, Exporting the configuration of SSB* (p. 97).

Always read the release notes of the firmware before updating SSB, because the release notes may include special instructions specific to the firmware version. [The release notes are available here.](#)

Steps:

- Step 1. Visit the BalaBit website and [download the latest firmware here.](#)
- Step 2. Navigate to the **Basic Settings > System** page.
- Step 3. To update the internal (core) firmware, select **Core firmwares**.
To update the external (boot) firmware, select **Boot firmwares**.



Main menu

- Basic Settings
 - Network
 - System**
 - High Availability
 - Date & Time
 - Management
 - Alerting & Monitoring
 - Troubleshooting
 - Dashboard
- AAA
- Policies
- Log
- Search
- Reports

User menu

- Private keystore
- Change password
- Preferences
- Logout

User info

User: admin
Host: 10.50.0.2
Last login: 2009-12-04 10:56 from 10.50.0.2

System monitor

Time: 2009-12-03 16:43
Remaining time: 09:29
HA: status: HA active: 08:00:27:1d:eb:cb
Modules: syslog-ng: **RUNNING**
Active Hosts: **UNKNOWN**
Senders: **UNKNOWN**
Load 1: 0.00 Load 15: 0.02
CPU Mem Disk Swap
100% ☐ ☐ ☐ ☐

System | Network | High Availability | Date & Time | Management | Alerting & Monitoring | Troubleshooting | Dashboard

System control

Traffic control

Version details

Export configuration

Import configuration

License

Sealed mode

Core firmwares

Upload new firmware: Choose... **Upload**

Available firmwares:

Slot	Version	Current	After reboot
1	HEAD1259758489	✓	✓ <input type="checkbox"/>
2			
3			
4			
5			

Boot firmwares

Upload new firmware: Choose... **Upload**

Available firmwares:

Slot	Version	Current	After reboot
1	HEAD1259758326	✓	✓ <input type="checkbox"/>
2			
3			
4			
5			

Figure 6.6. Managing the firmwares

For details on the different firmwares, see Section 2.7, *Firmware in SSB* (p. 9).

Step 4. Select the firmware file using the **Browse** button. The extension of firmware files is **.bin**

Step 5. Click **Upload**. After uploading, the new firmware is added to the **Available firmwares** list.

Step 6. Click ☐ in the **After reboot** column of the new firmware.

Step 7. Navigate to **System Control > This node** and click **Reboot**.

**Tip**

When SSB boots, it sends a message into the system log that includes the version numbers of both booted firmwares.

**Note**

If you experience any problems on the syslog-ng Store Box web interface after performing the upgrade, first empty the cache of your browser, or click the **Reload** button of your browser while holding the **Shift** key.



6.3.2. Procedure – Upgrading both the core and the boot firmware of a high availability system

Purpose:

If an SSB release requires the upgrading of both the boot firmware and the core (internal) firmware, complete the following steps:

Steps:

Step 1. Download both the core (internal) and the boot (external) firmware.

Step 2. Update the core firmware of the SSB using the web interface.

Step a. Navigate to **Basic Settings > System > Core firmwares** and upload the new core firmware.

Step b. When the upload is finished, select the **After reboot** option for the new firmware.



Warning

DO NOT REBOOT SSB AFTER UPGRADING THE CORE FIRMWARE.

Step 3. Repeat the previous step with the Boot firmware.

Step 4. Select **Basic Settings > High Availability > Reboot Cluster** to restart both nodes.



Warning

Hazard of data loss! As this step terminates all active connections, perform it only during maintenance hours to prevent data loss.





Note

If you experience any problems on the syslog-ng Store Box web interface after performing the upgrade, empty the cache of your browser, or click the Reload button of your browser while holding the Shift key.

6.3.3. Procedure – Reverting to an older firmware version

Purpose:

SSB can store up to five different firmware versions, any of them can be booted if required. The available firmwares are displayed on the **Basic Settings > System > Boot firmware** and **Basic Settings > System > Core firmware** pages. The list shows the detailed version of each firmware, including the version number, the revision number, and the build date. The firmware running on SSB is marked with  in the **Current** column. The firmware that will be run after the next SSB reboot is marked with  in the **After reboot** column.

To boot an older firmware, complete the following steps:

**Warning**

When upgrading SSB, it is possible that the configuration file is updated as well. In such cases, simply rebooting with the old firmware will not result in a complete downgrade, because the old firmware may not be able to read the new configuration file. If this happens, access the console menu of SSB, and select the **Revert Configuration** option to restore the configuration file to its state before the firmware was upgraded. For details on using the console menu, see *Section 6.4.1, Using the console menu of SSB (p. 100)*.

**Warning**

Downgrading from a feature release to an earlier (and thus unsupported) feature release, or to the stable release is not supported: this means that once you upgrade a system from a stable release (for example 1.0) to a feature release (for example 1.1), you will have to keep upgrading to the new feature releases until the next stable version release (for example 2.0) is published, or risk using an unsupported product.

Steps:

- Step 1. Navigate to **Basic Settings > System > Boot firmware**.
- Step 2. Select the firmware version to use, and click ☐ in the **After reboot** column.
- Step 3. Navigate to **Basic Settings > System > Core firmware**.
- Step 4. Select the firmware version to use, and click ☐ in the **Boot** column.
- Step 5. Select **System Control > This node > Reboot** to reboot SSB.
If you are running an SSB cluster, select **Basic Settings > High Availability > Reboot Cluster**.

6.3.4. Procedure – Updating the SSB license**Purpose:**

The SSB license must be updated before the existing license expires or when you purchase a new license. Information of the current license of SSB is displayed on the **Basic Settings > System > License** page. The following information is displayed:

Figure 6.7. Updating the license

- **Customer:** The company permitted to use the license (for example *Example Ltd.*).



- **Serial:** The unique serial number of the license.
- **Host limit:** The number of peers SSB accepts log messages from.
- **Validity:** The period in which the license is valid. The dates are displayed in YYYY/MM/DD format.

SSB gives an automatic alert one week before the license expires. An alert is sent also when the number of peers exceeds 90% of the limit set in the license.

To update the license, complete the following steps:



Warning

Before uploading a new license, you are recommended to backup the configuration of SSB. For details, see *Procedure 6.3.5, Exporting the configuration of SSB (p. 97)*.

Steps:

- Step 1. Navigate to **Basic Settings > System > License**.
- Step 2. Click **Browse** and select the new license file.
- Step 3. Click **Upload**, then **Commit**.
- Step 4. To activate the new license, navigate to **Traffic control > Syslog traffic** and click **Restart syslog-ng**.

6.3.5. Procedure – Exporting the configuration of SSB

Purpose:

The configuration of SSB can be exported (for manual archiving, or to migrate it to another SSB unit) from the **Basic Settings > System** page. Use the respective action buttons to perform the desired operation.

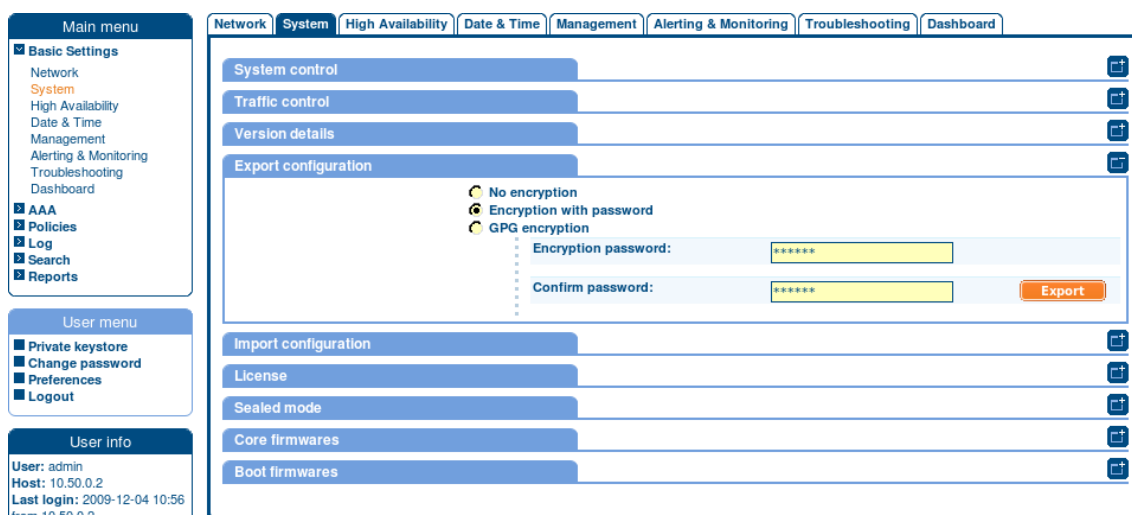


Figure 6.8. Exporting the SSB configuration

Steps:

- Step 1. Navigate to **Basic Settings > System > Export configuration**.



Step 2. Select how to encrypt the configuration:

- To export the configuration file without encryption, select **No encryption**.

**Warning**

Exporting the SSB configuration without encryption is not recommended, as it contains sensitive information such as password hashes and private keys.

- To encrypt the configuration file with a simple password, select **Encrypt with password** and enter the password into the **Encryption password** and **Confirm password** fields.
- To encrypt the configuration file with GPG, select **GPG encryption**. Note that this option uses the same GPG key that is used to encrypt automatic system backups, and is only available if you have uploaded the public part of a GPG key to SSB at **Basic Settings > Management > System backup**. For details, see *Procedure 4.7.3, Encrypting configuration backups with GPG* (p. 62).

Step 3. Click **Export**.

**Note**

The exported file is a gzip-compressed archive. On Windows platforms, it can be decompressed with common archive managers such as the *free 7-Zip tool*.

The name of the exported file is `<hostname_of_SSB>-YYYYMMDDTHHMM.config`; the `-encrypted` or `-gpg` suffix is added for password-encrypted and GPG-encrypted files, respectively.

6.3.6. Procedure – Importing the configuration of SSB

Purpose:

The configuration of SSB can be imported from the **Basic Settings > System** page. Use the respective action buttons to perform the desired operation.

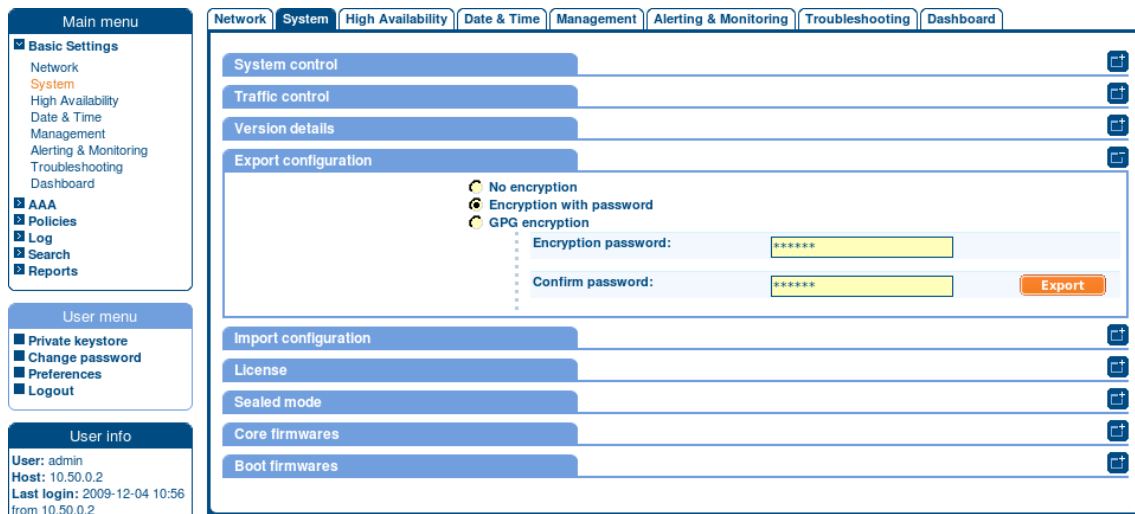


Figure 6.9. Importing the SSB configuration

**Warning**

It is possible to import a configuration exported from SSB 2.0 or 3.0 into SSB 3 LTS, but it is not possible to restore an 1.1 or 1.0 backup into 3 LTS.

Steps:

- Step 1. Navigate to **Basic Settings > System > Import configuration**.
- Step 2. Click **Browse** and select the configuration file to import.
- Step 3. Enter the password into the **Encryption password** field and click **Upload**.

**Warning**

When importing an older configuration, it is possible that there are logspaces on SSB that were created after the backing up of the old configuration. In such case, the new logspaces are not lost, but are deactivated and not configured. To make them accessible again, you have to:

1. Navigate to **Log > Spaces** and configure the logspace. Filling the **Access Control** field is especially important, otherwise the messages stored in the logspace will not be available from the **Search > Logs** interface.
2. Adjust your log path settings on the **Log > Paths** page. Here you have to re-create the log path that was sending messages to the logspace.

6.4. Accessing the SSB console

This section describes how to use the console menu of SSB, how to enable remote SSH access to SSB, and how to change the root password from the web interface.

**Tip**

If you need to find the SSB appliance in the server room, navigate to **Basic Settings > System > Hardware information > Blink system identification lights** and click **On**. This will blink the LEDs of hard disk trays on the front of the SSB appliance in red. Note that this is available only for *syslog-ng Store Box SSB5000*, and *SSB10000*.

6.4.1. Using the console menu of SSB

Connecting to the syslog-ng Store Box locally or remotely using Secure Shell (SSH) allows you to access the console menu of SSB. The console menu provides access to the most basic configuration and management settings of SSB. It is mainly used for troubleshooting purposes; the primary interface of SSB is the web interface.

The console menu is accessible to the *root* user using the password set during completing the Welcome Wizard.

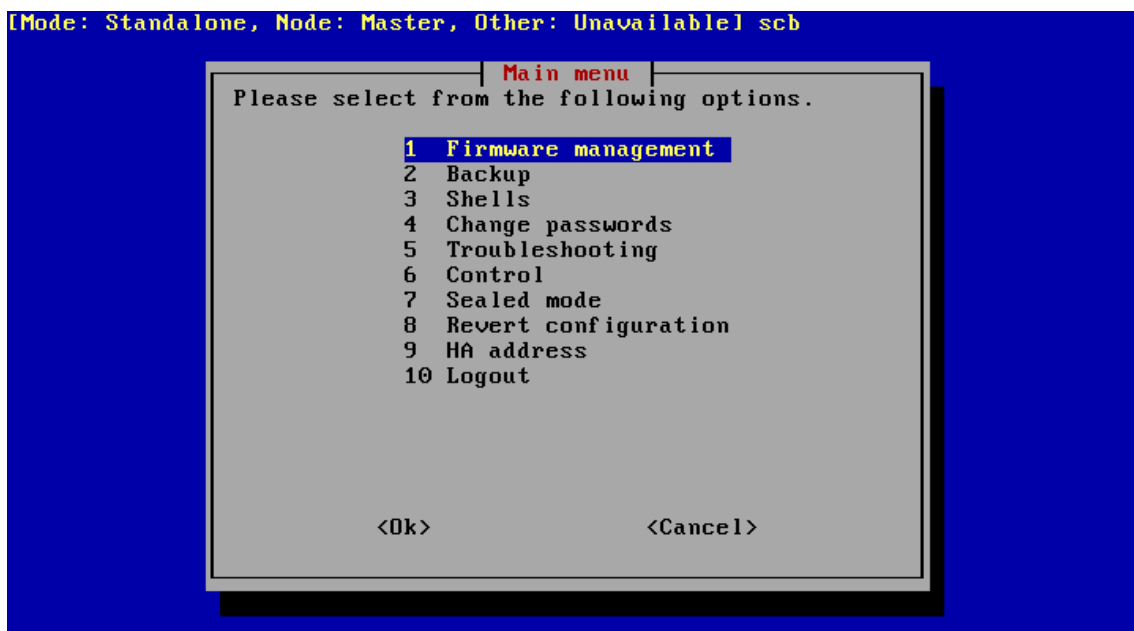


Figure 6.10. The console menu

The console menu provides allows you to perform the following actions:

- Select the active core and boot firmwares, and delete unneeded firmwares. Accessing the firmware management is useful if after an update the new firmware does not operate properly and the web interface is not available to activate the previous firmware.
- Start backup processes.
- Change the passwords of the *root* and *admin* users.
- Access the local shells of the core and boot firmwares. This is usually not recommended and only required in certain troubleshooting situations.
- Access the network-troubleshooting functions and display the available log files.
- Reboot and shutdown the system.
- Enable and disable sealed mode. For details, see *Section 6.7, Sealed mode (p. 103)*.



- Revert the configuration file. For details, see *Procedure 6.3.3, Reverting to an older firmware version (p. 95)*.
- Set the IP address of the HA interface.

**Note**

Note that logging in to the console menu automatically locks the SSB interface, meaning that users cannot access the web interface while the console menu is used. The console menu can be accessed only if there are no users accessing the web interface. The connection of web-interface users can be terminated to force access to the console menu.

6.5. Procedure – Enabling SSH access to the SSB host

Purpose:

Exclusively for troubleshooting purposes, you can access the SSB host using SSH. Completing the Welcome Wizard automatically disables SSH access. To enable it again, complete the following steps:

**Warning**

Accessing the SSB host directly using SSH is not recommended nor supported, except for troubleshooting purposes. In such case, the BalaBit Support Team will give you exact instructions on what to do to solve the problem.

Enabling the SSH server allows you to connect remotely to the SSB host and login using the *root* user. The password of the root user is the one you had to provide in the Welcome wizard. For details on how to change the root password from the web interface, see *Procedure 6.6, Changing the root password of SSB (p. 102)*

Steps:

Step 1. Navigate to **Basic Settings > Management > SSH settings**.



The screenshot shows the SSB Management interface with the 'Management' tab selected. On the left, there is a 'Main menu' with options like Basic Settings, Network, System, High Availability, Date & Time, Management (highlighted), Alerting & Monitoring, Troubleshooting, and Dashboard. Below this is a 'User menu' with Private keystore, Change password, Preferences, and Logout. Further down is 'User info' showing details for the 'admin' user. At the bottom left is a 'System monitor' section showing the current time and remaining session time. The main content area displays various system settings: Syslog, SNMP trap settings, SNMP agent settings, Mail settings, and SSH settings. The SSH settings section is expanded, showing 'Enable remote SSH access' (checked), 'Enable password authentication' (checked), and an 'Authorized keys' field with a plus icon. A 'Commit' button is located at the top right of the settings area.

Figure 6.11. Enabling remote SSH access to SSB

Step 2. Select the **Enable remote SSH access** option.



Note

Remote SSH access is automatically disabled if Sealed mode is enabled. For details, see *Section 6.7, Sealed mode (p. 103)*.

Step 3. Set the authentication method for the remote SSH connections.

- To enable password-based authentication, select the **Enable password authentication** option.
- To enable public-key authentication, click **+** in the **Authorized keys** field, click **+** and upload the private keys of the users who can access and manage SSB remotely via SSH.

Step 4. Click **Commit**.

The SSH server of SSB accepts connections only on the management interface if the management interface is configured. If the management interface is not configured, the SSH server accepts connections on the external interface. If possible, avoid enabling the SSH server of SSB when the management interface is not configured. For details on enabling the management connection, see *Procedure 4.3.1, Configuring the management interface (p. 39)*.

6.6. Procedure – Changing the root password of SSB

Purpose:



The root password is required to access SSB locally, or remotely via an SSH connection. Note that the password of the *root* user can be changed from the console menu as well. For details, see *Section 6.4, Accessing the SSB console (p. 99)*.

Steps:

Step 1. Navigate to **Basic Settings > Management > Change root password**.

Figure 6.12. Changing the root password of SSB

Step 2. Enter the new password into the **New root password** and **Confirm password** fields.



Note

SSB passwords can contain the following special characters: `!"#$%&'()*+,-./:;<=>?@[\\]^_`{|}`

Step 3. Click **Commit**.

6.7. Sealed mode

When sealed mode is enabled, the following settings are automatically applied:

- SSB cannot be accessed remotely via SSH for maintenance;
- the root password of SSB cannot be changed in sealed mode;
- Sealed mode can be disabled only from the local console. For details, see *Procedure 6.7.1, Disabling sealed mode (p. 104)*.

To enable sealed mode use one of the following methods:

- Select the **Sealed mode** option during the Welcome Wizard.
- Select **Basic Settings > System > Sealed mode > Activate sealed mode** on the SSB web interface.



- Login to SSB as root using SSH or the local console, and select **Sealed mode** > **Enable** from the console menu.

6.7.1. Procedure – Disabling sealed mode

Purpose:

To disable sealed mode, complete the following steps:

Steps:

- Step 1. Go to the SSB appliance and access the local console.
- Step 2. Login as *root*.
- Step 3. From the console menu, select **Sealed mode** > **Disable**
- Step 4. Select **Back to Main menu** > **Logout**.

6.8. Out-of-band management of SSB

SSB 3 LTS includes a dedicated out-of-band management interface conforming to the Intelligent Platform Management Interface (IPMI) v2.0 standards. The IPMI interface allows system administrators to monitor the system health of SSB and to manage the computer events remotely, independently of the operating system of SSB. SSB is accessible using the IPMI interface only if the IPMI interface is physically connected to the network.

- For details on connecting the IPMI interface, see *Procedure B.1, Installing the SSB hardware (p. 204)*.
- For details on configuring and using the IPMI interface to remotely monitor and manage SSB, see the following document:
 - The *Onboard BMC/IPMI User's Guide*, available at [the BalaBit Hardware Documentation page](#).

Basic information about the IPMI interface is available also on the SSB web interface on the **Basic Settings** > **High Availability** page. The following information is displayed:



Figure 6.13. Information about the IPMI interface SSB

- **Hardware serial number:** The unique serial number of the appliance.
- **IPMI IP address:** The IP address of the IPMI interface.
- **IPMI subnet mask:** The subnet mask of the IPMI interface.
- **IPMI default gateway IP:** The address of the default gateway configured for the IPMI interface.
- **IPMI IP address source:** Shows how the IPMI interface receives its IP address: dynamically from a DHCP server, or it uses a fixed static address.



Tip

If you need to find the SSB appliance in the server room, navigate to **Basic Settings > System > Hardware information > Blink system identification lights** and click **On**. This will blink the LEDs of hard disk trays on the front of the SSB appliance in red. Note that this is available only for *syslog-ng Store Box SSB5000*, and *SSB10000*.

6.9. Managing the certificates used on SSB

SSB uses a number of certificates for different tasks that can be managed from the **Basic Settings > Management > SSL certificate** menu.



Figure 6.14. Changing the web certificate of SSB

The following certificates can be modified here:

- **CA certificate:** The certificate of the internal Certificate Authority of SSB.
- **Server certificate:** The certificate of the SSB web interface, used to encrypt the communication between SSB and the administrators.

**Note**

If this certificate is changed, the browser of SSB users will display a warning stating that the certificate of the site has changed.

- **TSA certificate:** The certificate of the internal Timestamping Authority that provides the timestamps used when creating encrypted logstores.

For every certificate, the distinguished name (DN) of the X.509 certificate and the fingerprint of the private key is displayed. To display the entire certificate click on the DN; to display the public part of the private key, click on the fingerprint. It is not possible to download the private key itself from the SSB web interface, but the certificate can be downloaded in different formats (for example PEM, OpenSSH, Tectia).

**Note**

Other parts of SSB may use additional certificates that are not managed here.



During the initial configuration, SSB creates a self-signed CA certificate, and uses this CA to issue the certificate of the web interface (see **Server certificate**) and the internal Timestamping Authority (**TSA certificate**).

There are two methods to manage certificates of SSB:

- **Recommended:** Generate certificates using your own PKI solution and upload them to SSB. Generate a CA certificate and two other certificates signed with this CA using your PKI solution and upload them to SSB. For the Server and TSA certificates, upload the private key as well.

For details on uploading certificates and keys created with an external PKI, complete *Procedure 6.9.2, Uploading external certificates to SSB (p. 108)*.

**Warning**

The Server and the TSA certificates must be issued by the same Certificate Authority.

- Use the certificates generated on SSB. In case you want to generate new certificates and keys for SSB using its self-signed CA certificate, or generate a new self-signed CA certificate, complete *Procedure 6.9.1, Generating certificates for SSB (p. 107)*.

**Note**

Generate certificates using your own PKI solution and upload them to SSB whenever possible. Certificates generated on SSB cannot be revoked, and can become a security risk if they are somehow compromised.

6.9.1. Procedure – Generating certificates for SSB

Purpose:

Create a new certificate for the SSB webserver or the Timestamping Authority using the internal CA of SSB, or create a new, self-signed CA certificate for the internal Certificate Authority of SSB.

Steps:

Step 1. Navigate to **Basic Settings > Management > SSL certificate**.

Step 2. Fill the fields of the new certificate:

Step a. **Country:** Select the country where SSB is located (for example HU - Hungary).

Step b. **Locality:** The city where SSB is located (for example Budapest).

Step c. **Organization:** The company who owns SSB (for example Example Inc.).

Step d. **Organization unit:** The division of the company who owns SSB (for example IT Security Department).

Step e. **State or Province:** The state or province where SSB is located.

Step 3. Select the certificate you want to generate.



- To create a new certificate for the SSB web interface, select **Generate Server certificate**.
- To create a new certificate for the Timestamping Authority, select **Generate TSA certificate**.
- To create a new certificate for the internal Certificate Authority of SSB, select **Generate All**. Note that in this case new certificates are created automatically for the server and TSA certificates as well.

**Note**

When generating new certificates, the server and TSA certificates are signed using the certificate of the CA. If you have uploaded an external CA certificate along with its private key, it will be used to create the new server and TSA certificates. If you have uploaded an external CA certificate without its private key, use your external PKI solution to generate certificates and upload them to SSB.

**Warning**

Generating a new certificate automatically deletes the earlier certificate.

Step 4. Click **Commit**.

6.9.2. Procedure – Uploading external certificates to SSB

Purpose:

Upload a certificate generated by an external PKI system to SSB.

Prerequisites:

The certificate to upload. For the TSA and Server certificate, the private key of the certificate is needed as well. The certificates must meet the following requirements:

- SSB accepts certificates in PEM format. The DER format is currently not supported.
- SSB accepts private keys in PEM (RSA and DSA), PUTTY, and SSHCOM/Tectia format. Password-protected private keys are also supported.


For the internal CA certificate of SSB, uploading the private key is not required.

- For the TSA certificate, the *X509v3 Extended Key Usage* attribute must be enabled and set to *critical*. Also, its default value must be set to *Time Stamping*.
- For the Server certificate, the *X509v3 Extended Key Usage* attribute must be enabled and its default value set to *TLS Web Server Authentication*. Also, the Common Name of the certificate must contain the domain name or the IP address of the SSB host.

Steps:

Step 1. Navigate to **Basic Settings > Management > SSL certificate**.



Step 2. To upload a new certificate, click  next to the certificate you want to modify. A popup window is displayed.

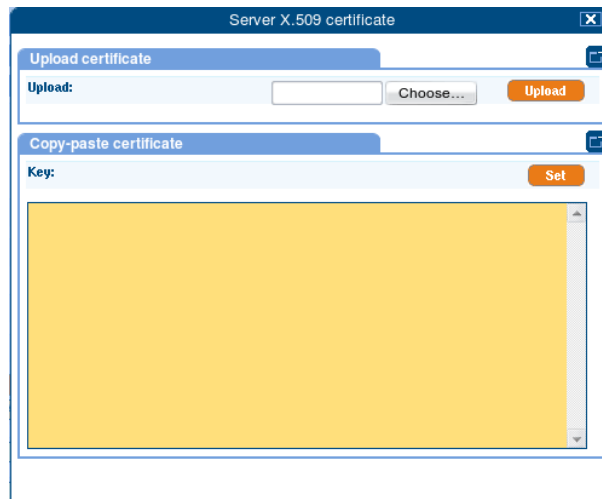



Figure 6.15. Uploading certificates

Select **Browse**, select the file containing the certificate, and click **Upload**. To upload a certificate chain, copy the certificates after each other in a single file. Alternatively, you can also copy-paste the certificate into the **Certificate** field and click **Set**. To copy-paste a certificate chain, copy and paste the certificates one by one after each other. The certificates do not have to be in order, SSB will order them. The chain is validated, if a member of the chain is missing, and error message is displayed.

Step 3. To upload the private key corresponding to the certificate, click the  icon next to the private key you want to modify. A popup window is displayed.

Select **Browse**, select the file containing the certificate, and click **Upload**. Alternatively, you can also copy-paste the certificate into the **Key** field and click **Set**.



Note

In case of a certificate chain, the private key has to be the same as the bottom level certificate.

Expected result:

The new certificate is uploaded. If you receive the *Certificate issuer mismatch* error message after importing a certificate, you must import the CA certificate which signed the certificate as well (the private key of the CA certificate is not mandatory).



Note

To download previously uploaded certificates, click on the certificate and either download the certificate (or certificate chain) in one single PEM or DER file, or you can download single certificate files separately (if it is a certificate chain).



6.10. Creating hostlist policies

SSB can use a list of host and network addresses at a number of places, for example for limiting the client that can send log messages to a log source, or the hosts that can access shared log spaces.

- For details on how to create a new hostlist, see *Procedure 6.10.1, Creating hostlists (p. 110)*.
- For details on how to import hostlists from a file, see *Procedure 6.10.2, Importing hostlists from files (p. 111)*.

6.10.1. Procedure – Creating hostlists

Purpose:

To create a new hostlist, complete the following steps.

Steps:

- Step 1. Navigate to **Policies > Hostlists** and select **+**.
- Step 2. Enter a name for the hostlist (for example *servers*).

Figure 6.16. Creating hostlists

- Step 3. Enter the IP address of the permitted host into the **Match > Address** field. You can also enter a network address in the *IP address/netmask* format (for example *192.168.1.0/24*). To add more addresses, click **+** and repeat this step.
- Step 4. To add hosts that are excluded from the list, enter the IP address of the denied host into the **Ignore > Address** field.

**Tip**

To add every address except for a few specific hosts or networks to the list, add the `0.0.0.0/0` network to the **Match** list, and the denied hosts or networks to the **Ignore** list.

Step 5. Click **Commit**.

**Warning**

If you modify a hostlist, navigate to **Basic Settings > System > Traffic control** and select **Restart syslog-ng** for the changes to take effect.

6.10.2. Procedure – Importing hostlists from files

Purpose:

To import hostlists from a text file, complete the following steps.

Steps:

Step 1. Create a plain text file containing the hostlist policies and IP addresses to import. Every line of the file will add an IP address or network to a policy. Use the following format:

name_of_the_policy;match or ignore;IP address

For example, a policy that ignores the `192.168.5.5` IP address and another one that matches on the `10.70.0.0/24` subnet, use:

```
policy1;ignore;192.168.5.5
policy2;match;10.70.0.0/24
```

To add multiple addresses or subnets to the same policy, list every address or subnet in a separate line, for example:

```
policy1;ignore;192.168.7.5
policy1;ignore;192.168.5.5
policy1;match;10.70.0.0/24
```

Step 2. Navigate to **Policies > Hostlists > Import from file > Browse** and select the text file containing the hostlist policies to import.



Figure 6.17. Importing hostlists

Step 3. If you are updating existing policies and want to add new addresses to them, select **Append**.
If you are updating existing policies and want to replace the existing addresses with the ones in the text file, select **Replace**.

Step 4. Click **Upload**, then **Commit**.

**Warning**

If you modify a hostlist, navigate to **Basic Settings > System > Traffic control** and select **Restart syslog-ng** for the changes to take effect.

6.11. Managing SAN access in SSB

SAN volumes can be managed using their own dedicated applications, for example the Sun Common Array Manager (CAM) in case of the Sun StorageTek modules. If you create new volumes on your SAN module, select **Basic Settings > Storage > Rescan** to make these volumes available from SSB.

**Warning**

Ensure that the LUN numbers assigned to the volume mappings of the storage are between 1 and 30 (including 1 and 30). Do NOT use the 0 or the 31 LUN, because SSB will not be able to access the volumes.

**Warning**

Do not unmap a volume that is used by SSB, as it can cause SSB to reboot.



Chapter 7. Configuring message sources

SSB receives log messages from remote hosts via *sources*. A number of sources are available by default, but you can also create new sources. Apart from the syslog protocols, SSB can also receive messages via the SNMP protocol, and convert these messages to syslog messages.

- For details on using the built-in message sources of SSB, see *Section 7.1, Default message sources in SSB (p. 114)*.
- For details on receiving SNMP messages, see *Procedure 7.2, Receiving SNMP messages (p. 114)*.
- For details on how to create new message sources, see *Procedure 7.3, Creating message sources in SSB (p. 116)*.

7.1. Default message sources in SSB

SSB automatically accepts messages from the following built-in sources:

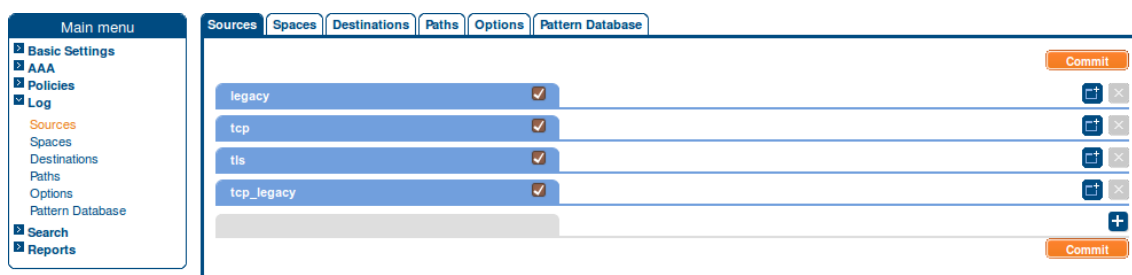


Figure 7.1. Default message sources in SSB

- **legacy**: Accepts UDP messages using the legacy BSD-syslog protocol on the port 514.
- **tcp**: Accepts TCP messages using the IETF-syslog protocol (RFC 5424) on port 601.
- **tls**: Accepts TLS-encrypted messages using the IETF-syslog protocol on port 6514. Mutual authentication is required: the client must show a (not necessarily valid) certificate; SSB sends the certificate created with the Welcome Wizard.
- **tcp_legacy**: Accepts TCP messages using the BSD-syslog protocol (RFC 3164) on port 514.

For the details of the various settings, see *Procedure 7.3, Creating message sources in SSB (p. 116)*.



Note
All default sources have name resolution enabled.

7.2. Procedure – Receiving SNMP messages

Purpose:



SSB can receive SNMP messages using the SNMPv2c protocol and convert these messages to syslog messages. SNMP messages are received using a special SNMP source that can be used in log paths like any other source. To configure receiving SNMP messages, complete the following steps:

Steps:

- Step 1. Navigate to **Log > Options > SNMP source**.
- Step 2. Ensure that the **SNMP source** option is enabled.

The screenshot shows the SSB configuration interface. On the left, there is a 'Main menu' with options like Basic Settings, AAA, Policies, Log, Sources, Spaces, Destinations, Paths, Options (highlighted), Pattern Database, Search, and Reports. Below it is a 'User menu' with Private key store, Change password, Preferences, and Logout. At the bottom left is 'User info' showing 'User: admin', 'Host: 10.40.255.254', and 'Last login: 2011-10-19 13:12 from 10.40.255.254'. The main content area has tabs for Sources, Spaces, Destinations, Paths, Options (selected), and Pattern Database. Under the 'Options' tab, the 'SNMP source' option is checked. Other settings include 'Use DNS' (checked), 'Community' (public), 'Hostlist' (a dropdown menu), 'Rate limit' (a text input field), 'Alerting' (checked), 'Artificial ignorance' (unchecked), 'Name resolving' (unchecked), 'Dashboard statistics' (checked), and 'Message rate alerting statistics' (unchecked). There are 'Commit' buttons at the top right and bottom right of the main content area.

Figure 7.2. Receiving SNMP messages

- Step 3. The default community of the SNMP messages is *public*. Modify the **Community** field if your hosts use a different community.



Note

SSB can receive messages only from a single community.

- Step 4. To limit which hosts can send SNMP messages to SSB, create a hostlist policy, add the permitted hosts to the policy, and select the policy from the **Hostlist** field. For details on creating hostlists, see *Section 6.10, Creating hostlist policies (p. 110)*.
- Step 5. To limit the rate of messages a host can send to SSB, enter the maximum number of packets (not messages) that SSB is allowed to accept from a single host into the **Rate limit** field. (This parameter sets the *hashlimit* parameter of the iptables packet filter that is applied to the source.)



Warning

When rate limiting is enabled, and a host sends a large number of messages, SSB processes only the amount set in the **Rate limit** field. Any additional messages are dropped, and most probably lost.



Step 6. To use name resolution for SNMP messages, enable the **Use DNS** option.

Step 7. Click **Commit**.

7.3. Procedure – Creating message sources in SSB

Purpose:

To create a custom message source, complete the following steps.

Steps:

Step 1. Navigate to **Log > Sources** and click **+**.

Step 2. Enter a name for the source into the top field. Use descriptive names that help you to identify the source easily.

Figure 7.3. Creating new message sources

Step 3. Select the interface of IP alias where SSB will receive the messages from the **Listening address** field.

Step 4. Enter the port number where SSB should accept the messages (for example 1999).



Step 5. If the information sent by the hosts to this source can be trusted, enable the **Trusted** option. SSB keeps the timestamps and the hostname of the messages sent by trusted clients. This corresponds to enabling the `keep_timestamp()` and `keep_hostname()` syslog-ng options for the source.

Step 6. In the **Transport** field, select the networking protocol (*UDP*, *TCP*, or *TLS*) that your clients use to transfer the messages to SSB.

When using TCP or TLS, you can set the maximum number of parallel connections in the **Maximum connections** field. This option corresponds to the `max_connections()` syslog-ng parameter.

When using TLS, SSB displays a certificate to the client. This certificate can be set at **Log > Options > TLS settings** (for details, see *Procedure 11.4, Setting the certificates used in TLS-encrypted log transport* (p. 153)). Optionally, SSB can perform mutual authentication and request and verify the certificate of the remote host (peer). Select the verification method to use from the **Peer verification** field.

- **None:** Do not request a certificate from the remote host, and accept any certificate if the host sends one.
- **Optional trusted:** If the remote host sends a certificate, SSB checks if it is valid (not expired) and that the Common Name of the certificate contains the domain name or the IP address of the host. If these checks fail, SSB rejects the connection. However, SSB accepts the connection if the host does not send a certificate.
- **Optional untrusted:** Accept any certificate shown by the remote host. Note that the host must show a certificate.
- **Required trusted:** Verify the certificate of the remote host. Only valid certificates signed by a trusted certificate authority are accepted. See *Procedure 6.9.2, Uploading external certificates to SSB* (p. 108) for details on importing CA certificates. Note that the Common Name of the certificate must contain the domain name or the IP address of the host.
- **Required untrusted:** SSB requests a certificate from the remote host, and rejects the connection if no certificate is received; if the certificate is not valid (expired); or if the Common Name of the certificate does not contain the domain name or the IP address of the host.



Warning

UDP is highly unreliable protocol, when using UDP, a large number of messages may be lost without any warning. Use TCP or TLS whenever possible.

Step 7. Select the syslog protocol used by the clients from the **Syslog protocol** field.

- If the clients use the legacy BSD-syslog protocol (RFC3164), select **Legacy**. This protocol is supported by most devices and applications capable to send syslog messages.
- If the clients use the new IETF-syslog protocol (for example the clients are syslog-ng 3.0 applications that use the `syslog` driver, or other drivers with the `flags(syslog-protocol)` option), select **Syslog**.

Step 8. Set the character **Encoding** and **Timezone** options of the incoming messages if needed.

Step 9. Select the **Use FQDN** option if you wish to store the full domain name of the sender host.



Step 10. Select the name resolving method to use from the **Use DNS** field.

Step 11. To accept messages only from selected hosts, create a hostlist and select it in the **Hostlist** field. For details on creating hostlists, see *Section 6.10, Creating hostlist policies (p. 110)*.

Step 12. If the messages arriving to the source do not comply to the standard syslog message format for some reason, select the **Do not parse messages** option. This option completely disables syslog message parsing and treats the complete as the MESSAGE part of a syslog message. Other information (timestamp, host, and so on) is added automatically by SSB.

If you still want to parse messages that comply to the standard syslog message format, but disable parsing for those that do not, select the **Ignore ambiguous program field** option. This will prevent SSB from treating the first word of the log message as the program name in case of non-standard syslog messages and thus resulting in unexpected behavior, for example polluting the statistics.

Step 13. To configure message rate alerting for the source, see *Procedure 4.6.4, Configuring message rate alerting (p. 50)*.

Step 14. Click **Commit**.

**Note**

Note that in order to actually store the messages arriving to this source, you have to include this source in a log path. For details, see *Chapter 10, Managing log paths (p. 144)*.



Chapter 8. Storing messages on SSB

SSB stores log messages in binary or plain-text log files called (log) spaces. These local destinations correspond to the `logstore()` and `file()` destinations of syslog-ng. Log spaces are stored locally on the hard disk of SSB, or on a SAN volume.



Note

Only the following edition of SSB can be connected to SAN devices: syslog-ng Store Box SANConnect.

- For details on which logspaces are created by default, see *Section 8.1, Default logspaces in SSB (p. 119)*.
- For notes and other important information on using encrypted log files (logstores), see *Section 8.3, Using logstores (p. 120)*.
- For details on how to create additional log spaces, see *Section 8.4, Creating custom message spaces in SSB (p. 121)*.
- For details on managing logspaces, see *Section 8.5, Managing log spaces (p. 127)*.
- For details on how to make the log files accessible remotely as a network drive, see *Section 8.6, Accessing log files across the network (p. 128)*.

8.1. Default logspaces in SSB

SSB has the following log spaces by default. Any incoming message is stored in these logspaces.



Warning

The default logspaces are not available if the SAN storage support is enabled in SSB. The following edition of SSB has SAN storage support: syslog-ng Store Box SANConnect. This edition of SSB can store log messages on the SAN volumes, and not directly on the hard disk of SSB.




Figure 8.1. Default logspaces in SSB

- **local**: An unencrypted, binary logspace for storing the log messages of SSB.
- **center**: An unencrypted, binary logspace for storing the log messages sent by the clients.



8.2. Configuring the indexer

Navigate to **Logs > Spaces** and select the desired logspace.

The indexer saves the indexes for the fields that are selected and makes them searchable. Indexing fields consumes disk space and processing power.

Enter the maximum amount of memory the indexer can use for the current logspace in the **Memory limit** field.

Select the desired fields to be indexed in the **Indexed fields**. The following fields can be indexed: **Facility**, **Priority**, **Program**, **Pid**, **Host**, **Tags**, **Name/value pairs**, **Message**. For the **Name/value pairs** field, select **All** to index all Name/value fields or enter the names to be indexed in the **Only with the name** field as comma-separated names. If the indexing of the **Message** field is enabled, the current **Delimiters** are displayed. By default, all indexers are selected.



Note

At least one field must be selected.



Note

It is not possible to search for whitespace () character in the MESSAGE part of the log message, since it is a hard-coded delimiter character.

8.3. Using logstores

This section contains important information about using logstore files for storing log messages, and describes the current limitations of the technology. These limitations will be addressed in future versions of SSB.

- In SSB version 1.0.x, it was not possible to browse the log messages stored in encrypted logstores from the SSB web interface. This problem has been addressed in SSB 1.1; for details, see *Section 12.3.5, Browsing encrypted log spaces (p. 171)*.
- Indexing logstore files is currently limited. The indexer can handle only one file from a logstore for every day (SSB automatically starts a new log file for every day, this corresponds to using the *\$DAY* macro of syslog-ng). However, if you use a filename template that separates log messages based on the sender host or application, or if you use a custom template that uses a finer time-based macro (for example *\$HOUR*), then currently only the first file for the day is indexed.
- Logstore files consist of chunks. In rare cases, if the syslog-ng application running on SSB crashes for some reason, it is possible that a chunk becomes broken: it contains log messages, but the chunk was not finished completely. However, starting with SSB version 2 F1 the syslog-ng application running on SSB processes log messages into a journal file before writing them to the logstore file. That way logstore files are consistent even during unexpected crash, avoiding losing messages. Similarly, if the indexer application crashes for some reason, it may be possible that some parts of a logstore file are not indexed, and therefore the messages from this part of the file do not appear in search results. This does not mean that the messages are lost. Currently it is not possible to reindex a file.



These limitations will be addressed in future versions of SSB.

8.3.1. Viewing encrypted logs with logcat

To access logstore files, you can either:

- access the logstores using a network share —: for details, see *Section 8.6, Accessing log files across the network (p. 128)*(recommended), or
- login to SSB locally or remotely using SSH.

To display the contents of a logstore file, use the `logcat` command supplied with `syslog-ng`, for example `logcat /var/log/messages.lgs`. To display the contents of encrypted log files, specify the private key of the certificate used to encrypt the file, for example `logcat -k private.key /var/log/messages.lgs`. The contents of the file are sent to the standard output, so it is possible to use `grep` and other tools to find particular log messages, for example `logcat /var/log/messages.lgs |grep 192.168.1.1`.

Every record that is stored in the logstore has a unique record ID. The `logcat` application can quickly jump to a specified record using the `-- seek` option.

For files that are in use by `syslog-ng`, the last chunk that is open cannot be read. Chunks are closed when their size reaches the limit set in the `chunk_size` parameter, or when the time limit set in the `chunk_time` parameter expires and no new message arrives.

When the logstore file is encrypted, a hash is also generated for every chunk to verify the integrity of the chunk. The hashes of the chunks are chained together to prevent injecting chunks into the logstore file. The encryption algorithm used is `aes128` in CBC mode, the hashing (HMAC) algorithm is `hmac-sha1`.



Warning

If the `syslog-ng` Premium Edition application or the computer crashes, an unclosed chunk remains at the end of the file. This chunk is marked as broken, its data stays there but is not shown by `logcat`.

8.4. Creating custom message spaces in SSB

To create a custom log space, complete one of the following procedures:

- Store the log messages in binary logstore files, complete *Procedure 8.4.1, Creating a new logstore (p. 121)*.
- Store the log messages in traditional plain-text files, complete *Procedure 8.4.2, Creating a new text logspace (p. 124)*.

8.4.1. Procedure – Creating a new logstore

Steps:

- Step 1. Navigate to **Log > Spaces** and click **+**.
- Step 2. Enter a name for the log space into the top field. Use descriptive names that help you to identify the source easily. Note that the name of the logspace must begin with a number or a letter.



The screenshot shows the SSB web interface with the 'Spaces' tab active. On the left, there's a 'Main menu' with options like Basic Settings, AAA, Policies, Log, Sources, Spaces, Destinations, Paths, Options, Pattern Database, Search, and Reports. Below it is a 'User menu' with Private keystore, Change password, Preferences, and Logout. A 'User info' section shows the user is 'admin' with host '10.40.255.254' and last login on '2012-03-23 15:18'. A 'System monitor' section displays system status: Time: 2012-03-23 16:37, Remaining time: 08:06, Locked: admin@10.40.255.254, Modules: syslog-ng: RUNNING, Active: Hosts: 2, Senders: 2, Load 1: 0.19, Load 15: 0.26, and CPU/Mem/Disk/Swap usage bars (1%, 61%, 9%, 0%).

The main configuration area has tabs for Sources, Spaces, Destinations, Paths, Options, and Pattern Database. Under 'Spaces', there are buttons for Backup ALL, Restore ALL, Archive/cleanup ALL, and Empty ALL, and a 'Commit' button. A list of spaces shows 'local', 'center', and 'newlogstore' (highlighted). The 'Type' field is set to 'LogStore'. The 'Compressed logstore' checkbox is checked. The 'Encryption certificate' field has a certificate icon. The 'Decryption private keys' field has a plus icon. The 'Timestamping frequency' is set to 600 seconds. The 'Indexer' is set to 'Disabled'. The 'Filename template' is set to 'All messages in one file'. Other fields include Backup policy, Archive/Cleanup policy, Sharing policy, Warning size (set to Gb), Access control (set to Group), and Message rate alerting (unchecked). A 'Commit' button is at the bottom right.

Figure 8.2. Creating a new logstore

- Step 3. Select **LogStore** from the **Type** field.
- Step 4. To encrypt the log files using public-key encryption, click in the **Encryption certificate** field. A popup window is displayed. Click **Browse**, select the certificate you want to use to encrypt the log files, then click **Upload**. Alternatively, you can paste the certificate into the **Certificate** field and click **Upload**.

**Note**

To view encrypted log messages, you will need the private key of this certificate. For details on browsing encrypted logstores online on the SSB web interface, see *Section 12.3.5, Browsing encrypted log spaces (p. 171)*. Encrypted log files can be displayed using the `logcat` command-line tool as well. The `logcat` application is currently available only for UNIX-based systems.

- Step 5. By default, SSB requests a timestamp every ten minutes from the internal Timestamping Authority. Adjust the frequency of timestamping requests in the **Timestamping frequency** field if needed. For



details on how to request timestamps from an external provider, see *Section 11.2, Timestamping configuration on SSB (p. 151)*.

Step 6. To automatically index the logstore files, select the **Enable** option of the **Indexer** field.

To limit the number of hits when searching in the logstore, enter the maximum number of search result hits in the **Maximum number of search results** field. To disable the limit, enter 0.

By default, the following fields are indexed, if indexing is enabled: **Program, Host, Name/value pairs, Message**.

By default, the indexer uses the following delimiter characters to separate the message into words (tokens): `:&~?![]=, ; () ' "`. If your messages contain segments that include one of these delimiters, and you want to search for these segments as a whole, remove the delimiter from the list. For example, if your log messages contain MAC addresses, and you want to be able to search for messages that contain a particular MAC address, delete the colon (:) character from the list of delimiters. Otherwise, the indexer will separate the MAC address into several tokens.

Step 7. Logstore files are compressed by default. If you do not want to use compression, uncheck the **Compressed logstore** option.

Step 8. Select how to organize the log files of this log space from the **Filename template** field.

- To save every message received during a day into a single file, select **All messages in one file**.
- To create a separate log file for every peer (IP address or hostname) that sends messages, select the **Per host** option. This option corresponds to using the `$HOST` macro of syslog-ng.
- To create a separate log file for every application that sends messages, select the **Per application** option. This option corresponds to using the `$PROGRAM` macro of syslog-ng.
- To create a separate log file for every application of every peer (IP address or hostname) that sends messages, select **Per host and application** option. This option corresponds to using the `$HOST-$PROGRAM` macros of syslog-ng.
- To specify a custom template for naming the log files, select the **Custom** option and enter the template into the appearing **Template** field. For details on using filename templates, see *The syslog-ng Premium Edition Administrator Guide*.

Step 9. To create automatic daily backups of the log space to a remote server, create a backup policy and select it from the **Backup policy** field. For details on creating backup policies, see *Section 4.7, Data and configuration archiving and backups (p. 54)*.

Step 10. To archive the log space automatically daily, create an archiving policy and select it from the **Archive/Cleanup policy** field. For details on creating archiving policies, see *Section 4.7, Data and configuration archiving and backups (p. 54)*.



Warning

Use archiving and cleanup policies to remove older logfiles from SSB, otherwise the hard disk of SSB may become full.



- Step 11. To make the log files of this log space available via the network, create a sharing policy and select it from the **Sharing policy** field. For details on creating sharing policies, see *Section 8.6, Accessing log files across the network (p. 128)*.
- Step 12. If you are using a SAN storage (available for syslog-ng Store Box SANConnect) select the storage volume that will store the messages. Note that specifying a volume is mandatory for these SSB editions, as they do not store logs on their local hard disk. Connections to SAN devices can be set on **Basic Settings > Storage**.

**Note**

When modifying an existing log space, assigning a new volume to the log space does neither delete nor copy the messages from the old volume.

- Step 13. Set a size for the log space in the **Warning size** field: SSB will send an alert if the size of this log space exceeds the limit.

**Warning**

Make sure that the **Logspace exceeded warning size** alert is enabled in **Basic Settings > Alerting & Monitoring** page, and that the mail and SNMP settings of the **Basic Settings > Management** page are correct. Otherwise, you will not receive any alert when the log space exceeds the size limit. For details on alerting and monitoring, see also *Section 4.6, Configuring system monitoring on SSB (p. 47)*.

- Step 14. By default, members of the *search* group can view the stored messages online. Use the **Access control** option to control which usergroups can access the log space. For details, see also *Section 5.6, Managing user rights and usergroups (p. 77)*.

- Step 15. Click **Commit**.

8.4.2. Procedure – Creating a new text logspace

Steps:

- Step 1. Navigate to **Log > Spaces** and click **+**.
- Step 2. Enter a name for the log space into the top field. Use descriptive names that help you to identify the source easily.



Figure 8.3. Creating a new text logspace

Step 3. Select **Text file** from the **Type** field.

Step 4. Select the template to use for the messages. The following templates are available:

- *Legacy*: `template("$DATE $HOST $MSGHDR$MSG\n")`
- *ISO date*: `template("$ISODATE $HOST $MSGHDR$MSG\n")`
- *Custom*: Specify a custom syslog-ng template in the appearing **Template** field. For details on using templates, see *The syslog-ng Premium Edition Administrator Guide*.

Step 5. Adjust the number of messages that are stored in the memory in the **Memory buffer size** field. This parameter corresponds to the `log_fifo_size()` parameter of syslog-ng.

Step 6. Select how to organize the log files of this log space from the **Filename template** field.

- To save every message received during a day into a single file, select **All messages in one file**.
- To create a separate log file for every peer (IP address or hostname) that sends messages, select the **Per host** option. This option corresponds to using the `$HOST` macro of syslog-ng.



- To create a separate log file for every application that sends messages, select the **Per application** option. This option corresponds to using the `$PROGRAM` macro of syslog-ng.
- To create a separate log file for every application of every peer (IP address or hostname) that sends messages, select **Per host and application** option. This option corresponds to using the `$HOST-$PROGRAM` macros of syslog-ng.
- To specify a custom template for naming the log files, select the **Custom** option and enter the template into the appearing **Template** field. For details on using filename templates, see *The syslog-ng Premium Edition Administrator Guide*.

Step 7. To create automatic daily backups of the log space to a remote server, create a backup policy and select it from the **Backup policy** field. For details on creating backup policies, see *Section 4.7, Data and configuration archiving and backups (p. 54)*.

Step 8. To archive the log space automatically daily, create an archiving policy and select it from the **Archive/Cleanup policy** field. For details on creating archiving policies, see *Section 4.7, Data and configuration archiving and backups (p. 54)*.



Warning

Use archiving and cleanup policies to remove older logfiles from SSB, otherwise the hard disk of SSB may become full.

Step 9. To make the log files of this log space available via the network, create a sharing policy and select it from the **Sharing policy** field. For details on creating sharing policies, see *Section 8.6, Accessing log files across the network (p. 128)*.

Step 10. If you are using a SAN storage (available for syslog-ng Store Box SANConnect) select the storage volume that will store the messages. Note that specifying a volume is mandatory for these SSB editions, as they do not store logs on their local hard disk. Connections to SAN devices can be set on **Basic Settings > Storage**.



Note

When modifying an existing log space, assigning a new volume to the log space does neither delete nor copy the messages from the old volume.

Step 11. Set a size for the log space in the **Warning size** field: SSB will send an alert if the size of this log space exceeds the limit.



Warning

Make sure that the **Logspace exceeded warning size** alert is enabled in **Basic Settings > Alerting & Monitoring** page, and that the mail and SNMP settings of the **Basic Settings > Management** page are correct. Otherwise, you will not receive any alert when the log space exceeds the size limit. For details on alerting and monitoring, see also *Section 4.6, Configuring system monitoring on SSB (p. 47)*.



Step 12. By default, members of the *search* group can view the stored messages online. Use the **Access control** option to control which usergroups can access the log space. For details, see also *Section 5.6, Managing user rights and usergroups (p. 77)*.

Step 13. Click **Commit**.

8.5. Managing log spaces

Log spaces are mostly managed automatically using backup and archiving policies, as described in *Section 4.7, Data and configuration archiving and backups (p. 54)*. However, backup and archiving can be started manually as well. To display the details of a log space, click . A number of action buttons is shown in the top row.

The screenshot shows the SSB web interface for managing log spaces. On the left is a sidebar with a 'Main menu' (Basic Settings, AAA, Policies, Log, Sources, Spaces, Destinations, Paths, Options, Pattern Database, Search, Reports) and a 'User menu' (Private keystore, Change password, Preferences, Logout). Below these is 'User info' (User: admin, Host: 10.50.0.6, Last login: 2010-06-17 11:35) and a 'System monitor' (Time: 2010-06-17 12:24, Remaining time: 09:40, Locked: admin@10.50.0.6, Modules: syslog-ng: STOPPED, Active Hosts: 0, Senders: 0, Load 1: 0, Load 15: 0, CPU: 100%, Mem: 0%, Disk: 21%, Swap: 0%).

The main content area has tabs for Sources, Spaces, Destinations, Paths, Options, and Pattern Database. The 'Spaces' tab is active, showing a list of log spaces. The 'local' space is selected, displaying its details. At the top of the details are buttons: Backup ALL, Restore ALL, Archive/cleanup ALL, Empty ALL, and a Commit button. Below these are buttons for Backup, Restore, Archive/Cleanup, and Empty. The 'Size' row shows '2.23 MB (last updated: 2010-06-17 12:10)' and a 'Get current size' button. The 'Type' section has radio buttons for 'LogStore' (selected) and 'Text file'. Under 'LogStore', there are fields for 'Compressed logstore' (checked), 'Encryption certificate' (empty), 'Decryption private keys' (empty), 'Timestamping frequency' (600 seconds), and 'Indexer' (radio buttons for Enabled and Disabled, with Disabled selected). The 'Filename template' section has radio buttons for 'All messages in one file' (selected), 'Per host', 'Per application', 'Per host and application', and 'Custom'. Below this are dropdown menus for 'Backup policy', 'Archive/Cleanup policy', and 'Sharing policy'. The 'Warning size' is set to 'Gb'. The 'Access control' section has a 'Group' dropdown menu with 'search' selected.

Figure 8.4. Managing log spaces



Tip

The size of the log space is displayed in the **Size** row of the log space details. To refresh the data, select **Get current size**.

- To start the backup process manually, click **Backup**.
- To restore the log files from the backup server to SSB click **Restore**.

**Warning**

Restoring the backup replaces every log file of the log space with the files from the backup. Any log message saved into the log space since the backup is irrevocably lost.

- To start the archiving and the cleanup process manually, click **Archive/Cleanup**.

**Warning**

If the archiving policy selected for the log space is set to perform only cleanup, log messages older than the Retention Time are deleted and irrevocably lost. For details, see *Section 4.7, Data and configuration archiving and backups (p. 54)*.

- To delete every log file in the log space, click **Empty**. This option can be useful if you have to quickly free up space on SSB, or if you want to delete a log space.

**Warning**

This action deletes every file of the log space. Any log message not archived or backed up is irrevocably lost.

Similar action buttons are available at the top of the **Log > Spaces** page to backup, archive, or delete the contents of every logspace. These actions are performed on every logspace with their respective settings, that is, clicking **Backup All** creates a backup of every logspace using the backup policy settings of the individual logspace.

8.6. Accessing log files across the network

The log files stored on SSB can be accessed as a network share if needed using the Samba (CIFS) or Network File System (NFS) protocols. Sharing is controlled using policies that specify the type of the share and the clients (hosts) and users who can access the log files. Sharing is possible also if SSB is part of a domain.

- If you manage SSB users locally, users who have SSB account can access the shared folders. Complete *Procedure 8.6.1, Sharing log files in standalone mode (p. 128)*.
- If you manage SSB users from LDAP, you must join SSB to your domain. Complete *Procedure 8.6.2, Sharing log files in domain mode (p. 130)*.
- For details on how to access the shared files, see *Section 8.6.3, Accessing shared files (p. 133)*.

8.6.1. Procedure – Sharing log files in standalone mode

Steps:

Step 1. Navigate to **Policies > Shares > SMB/CIFS options** and select **Standalone mode**.



Figure 8.5. Sharing log spaces

Step 2. Select **+** to create a new share policy and enter a name for the policy.

Step 3.

Figure 8.6. Creating share policies

Select the type of the network share from the **Type** field.

- To access the log files using NFS (Network File System), select **NFS**.
- To access the log files using Samba (Server Message Block protocol), select **CIFS**.

Step 4. If you are using the Samba protocol, you can control which users and hosts can access the shares. Otherwise, every user with an SSB account has access to every shared log file.

- To control which users can access the shared files, enter the name of the usergroup who can access the files into the **Allowed group** field. For details on local user groups, see *Procedure 5.3, Managing local usergroups (p. 72)*.
- To limit the hosts from where the shares can be accessed, create a hostlist and select it from the **Hostlist** field. For details on creating hostlists, see *Section 6.10, Creating hostlist policies (p. 110)*.



Step 5. Click **Commit**.

Step 6. To display the details of the log space, navigate to **Log > Spaces** and click

Step 7.

The screenshot displays the 'Log Space' configuration page for a space named 'center'. The interface includes a sidebar with navigation menus and a main configuration area. The 'center' log space is currently empty. The configuration options are as follows:

- Logspace is empty** (Restore button)
- Size:** 5 kB (last updated: 2010-06-17 13:26) (Get current size button)
- Type:**
 - ☒ LogStore
 - ☐ Text file
- Compressed logstore:** ☒
- Encryption certificate:** ☒
- Decryption private keys:**
- Timestamping frequency:** 600 seconds
- Indexer:**
 - ☐ Enabled
 - ☒ Disabled
- Filename template:**
 - ☒ All messages in one file
 - ☐ Per host
 - ☐ Per application
 - ☐ Per host and application
 - ☐ Custom
- Backup policy:**
- Archive/Cleanup policy:**
- Sharing policy:** myshare
- Warning size:** Gb
- Access control:**
 -
 -

Figure 8.7. Setting the share policy of a log space

Select the share policy to use from the **Sharing policy** field.

Step 8. Click **Commit**.

8.6.2. Procedure – Sharing log files in domain mode

Steps:

Step 1. Navigate to **Policies > Shares > SMB/CIFS options** and select **Domain mode**.

Step 2. Enter the name of the domain (for example *mydomain*) into the **Domain** field.



Figure 8.8. Joining a domain

- Step 3. Enter the name of the realm (for example *mydomain.example.com*) into the **Full domain name** field.

**Note**

Ensure that your DNS settings are correct and that the full domain name can be resolved from SSB. To check this, navigate to **Basic Settings > Troubleshooting > Ping**, enter the full domain name into the **Hostname** field, and select **Ping host**.

- Step 4. Click **Join domain**. A popup window is displayed.
- Step 5. SSB requires an account to your domain to be able to join the domain. Enter the name of the user into the **Username** field, and the corresponding password into the **Password** field. Optionally, you can enter the name of your domain controller into the **Domain controller** field. If you leave this field blank, SSB will try to find the domain controller automatically.

**Note**

Ensure that your DNS settings are correct and that the hostname of the domain controller can be resolved from SSB. To check this, navigate to **Basic Settings > Troubleshooting > Ping**, enter the name of the domain controller into the **Hostname** field, and select **Ping host**.

- Step 6. Click **Join domain**.
- Step 7. Select **+** to create a new share policy and enter a name for the policy.



Figure 8.9. Creating share policies

Step 8. Select the type of the network share from the **Type** field.

- To access the log files using NFS (Network File System), select **NFS**.
- To access the log files using Samba (Server Message Block protocol), select **CIFS**.

Step 9. If you are using the Samba protocol, you can control which users and hosts can access the shares. Otherwise, every user with an SSB account has access to every shared log file.

- To control which users can access the shared files, enter the name of the LDAP group who can access the files into the **Allowed group** field. Note that the users and SSB must be members of the same domain.
- To limit the hosts from where the shares can be accessed, create a hostlist and select it from the **Hostlist** field. For details on creating hostlists, see *Section 6.10, Creating hostlist policies (p. 110)*.

Step 10. Click **Commit**.

Step 11. To display the details of the log space, navigate to **Log > Spaces** and click



Step 12.

Figure 8.10. Setting the share policy of a log space

Select the share policy to use from the **Sharing policy** field.

Step 13. Click **Commit**.

8.6.3. Accessing shared files

This section describes how to access log files that are shared using a share policy. For details on sharing log files, see *Section 8.6, Accessing log files across the network (p. 128)*.

Every shared log space is available as a separate shared folder, even if they all use a single share policy. The name of the shared folder is the name of the log space. Within the shared folder, the log files are organized into the following directory structure: YEAR/MM-DD/. The files are named according to the filename template set for the log space; the extension of logstore files is *.store*, while the extension of text files is *.log*. Note that the root directory of the share may also contain various files related to the log space, like index files for logstores. All files are read-only.

**Note**

When using NFS for sharing the log space, the name of the shared folder will be the following: `/exports/{logspace_id}/...`. The following example demonstrates how to mount a shared log space using NFS on Linux.

**Example 8.1. Mounting a shared log space using NFS on Linux**

```
mount -t nfs {ssb_ip}:/exports/{logspace_id} /mnt/testmount
```



Chapter 9. Forwarding messages from SSB

SSB can forward log messages to remote destinations. The remote destination can be an SQL database running on a remote server, or a syslog or log analyzing application running on a remote server.


- To forward messages to a remote SQL database, complete *Procedure 9.1, Forwarding log messages to SQL databases (p. 135)*. Currently Oracle, Microsoft SQL (MSSQL), MySQL, and PostgreSQL databases are supported.
- To forward messages to a remote server, complete *Procedure 9.3, Forwarding log messages to remote servers (p. 139)*.

9.1. Procedure – Forwarding log messages to SQL databases

Purpose:

This section describes how to forward log messages from SSB to a remote SQL database server.

Steps:

Step 1. To create a new remote destination, navigate to **Log > Destinations** and select .

Step 2. Enter a name for the destination.



Note

This name will be used in the name of the database tables created by SSB. For compatibility reasons, it can contain only numbers, lowercase characters, and the underscore () character, for example *example_database_destination*.

Step 3. Select **Database Server**.



The screenshot shows the Balabit Syslog-ng configuration web interface. On the left, there is a sidebar with a 'Main menu' (Basic Settings, AAA, Policies, Log, Sources, Spaces, Destinations, Paths, Options, Pattern Database, Search, Reports), a 'User menu' (Private keystore, Change password, Preferences, Logout), and 'User info' (User: admin, Host: 10.40.255.254, Last login: 2011-11-07 12:03 from 10.40.255.254). Below this is a 'System monitor' section showing system time, remaining time, and resource usage (CPU, Mem, Disk, Swap) with bar charts.

The main content area is titled 'Destinations' and contains a form for creating a new destination. The 'Type' field is set to 'Database server'. The 'Database type' is 'PostgreSQL'. The 'Address' is 'database.example.com', 'Port' is '5432', 'Username' is 'ssb', and 'Password' is masked. The 'Database name' is 'logs'. The 'Test connection' button is visible. 'Flush lines' is 1000 and 'Flush timeout' is 500 milliseconds. 'Table rotation' is set to 'Monthly'. 'Table schema' is 'Extended'. 'Retention time' is 31 days. 'Access control' is set to 'search'. At the bottom, there are fields for 'Timestamp fractions of a second' (0 digits), 'Timezone', 'Output disk buffer' (0 Mb), 'Output memory buffer' (10000 messages), and 'Message rate alerting' (unchecked). A 'Commit' button is at the bottom right.

Figure 9.1. Creating database destinations

- Step 4. Select the type of the remote database from the **Database type** field.
- Step 5. Enter the IP address or hostname of the database server into the **Address** field. If the database is running on a non-standard port, adjust the **Port** setting.
- Step 6. Enter the name and password of the database user account used to access the database into the **Username** and **Password** fields, respectively. This user needs to have the appropriate privileges for creating new tables.
- Step 7. Enter the name of the database that will store the log messages into the **Database name** field.



Step 8. *Optional step:* Enter the number of log message lines into the **Flush lines** field that SSB should wait before sending them off in a single batch. Setting this number high increases throughput as fully filled frames are sent to the network. However, it also increases message latency. This latency can be limited by using the **Flush timeout** option.



Note

Flush lines is in connection with the **Output memory buffer** value. (To set the **Output memory buffer** value, navigate to **Log > Destinations**). The value of **Output memory buffer** has to be greater than or equal to the value of **Flush lines**.

- Step 9. *Optional step:* Enter the number of milliseconds into the **Flush timeout** field that SSB should wait before sending them off in a single batch if there has not been as many logs as specified in the **Flush lines** field. Setting this number high increases throughput as fully filled frames are sent to the network. However, it also increases message latency. This latency can be limited by using the **Flush lines** option.
- Step 10. SSB will automatically start a new table for every day or every month. Optionally, you can also create custom tables. Select the table naming template from the **Table rotation** field.
- Step 11. Select which columns should SSB insert into the database. You can use one of the predefined templates, or select **Custom columns** to create a custom template. The available templates are described in *Section 9.2, SQL templates in SSB (p. 138)*.
- Step 12. SSB can automatically delete older messages and tables from the database. By default, messages are deleted after one month. Adjust the **Retention time** as needed for your environment.
- Step 13. The logs stored in the database can be accessed using the search interface of SSB. Enter the name of the usergroup who can access the logs into the **Access control > Group** field. To add more groups (if needed), click **+**.
- Step 14. The timestamps of most log messages is accurate only to on second. SSB can include more accurate timestamps: set how many digits should be included in the **Timestamp fractions of a second** field. This option corresponds to the `frac_digits()` parameter of syslog-ng.
- Step 15. If the server and SSB are located in a different timezone and you use the *Legacy* message template (which does not include timezone information), select the timezone of the server from the **Timezone** field.
- Step 16. Set the size of the disk buffer in the **Output disk buffer** field. If the remote server becomes unavailable, SSB will buffer messages to the hard disk, and continue sending the messages when the remote server becomes available. This option corresponds to the `log_disk_fifo_size()` parameter of syslog-ng.
- Step 17. By default, SSB buffers up to 10000 messages in its memory if the remote server cannot accept them fast enough. To modify this value, adjust the **Output memory buffer** field as needed. This option corresponds to the `log_fifo_size()` parameter of syslog-ng.
- Step 18. Click **Commit**.
- Step 19. To start sending messages to the destination, include the new destination in a logpath. For details, see *Chapter 10, Managing log paths (p. 144)*.
- Step 20. To test if the database is accessible, select **Test connection**.



9.2. SQL templates in SSB

The following sections describe the SQL templates available in SSB:

- Legacy
- Full
- Custom

9.2.1. The Legacy template

The **Legacy** template stores messages in the `ssb_sql_messages_${R_YEAR}_${R_MONTH}` table. The following columns are created:

- `insert_time`: The date when SSB received the message in Unixtime format.
- `rule_id`:
- `__row_id`:
- `date_time`: The date the message was sent in `YEAR-MONTH-DAY HOUR:MINUTE:SECOND` format.
- `facility`: The facility that sent the message.
- `priority`: The priority level of the message.
- `host`: The IP address or hostname of the host where the message was generated.
- `program`: The name of the application that generated the message.
- `pid`: The ID number of the process that generated the message (this field is automatically set to zero if the PID is not included in the message).
- `message`: The text of the log message.

The `insert_time`, `rule_id`, `date_time`, `facility`, `host`, and `program` columns are indexed.

9.2.2. The Full template

The **Full** template stores messages in the `ssb_sql_messages_${R_YEAR}_${R_MONTH}` table. The following columns are created:

- `insert_time`: The date when SSB received the message in Unixtime format.
- `rule_id`:
- `__row_id`:
- `date_time`: The date the message was sent in `YEAR-MONTH-DAY HOUR:MINUTE:SECOND` format.
- `facility`: The facility that sent the message.
- `priority`: The priority level of the message.
- `sourceip`: The IP address of the host that sent the message.
- `host`: The IP address or hostname of the host where the message was generated.
- `program`: The name of the application that generated the message.



- *pid*: The ID number of the process that generated the message (this field is automatically set to zero if the PID is not included in the message).
- *message*: The text of the log message.

The *insert_time*, *rule_id*, *date_time*, *facility*, *host*, *sourceip*, and *program* columns are indexed.

9.2.3. The Custom template

The **Custom** template allows you to specify the columns to use. Enter a name for the column, select its type, and specify its content using macros. For details on using macros, see [*The syslog-ng Premium Edition Administrator Guide*](#). Select the **Indexed** option if you want the database to index the column.

9.3. Procedure – Forwarding log messages to remote servers

Purpose:

This section describes how to forward messages from SSB to a remote server.

Steps:

- Step 1. Navigate to **Log > Destinations** and select **+** to create a new remote destination.
- Step 2. Select **Remote host**.



Figure 9.2. Creating server destinations

Step 3. Enter the IP address or hostname of the remote server into the **Address** field. Enter the port where the server is accepting syslog messages into the **Port** field.

Step 4. Select the network protocol used to transfer the log messages from the **Transport** field. The UDP, TCP, and the encrypted TLS protocols are available. The UDP and TLS protocols have additional parameters.

When forwarding messages using UDP, the remote host will see the messages as if they originated from SSB. Select the **Spoof source address** option to make them seem to originate from their original sender.

For TLS, select a method to verify the identity of the remote host. The following options are available:

- *None*: Do not request a certificate from the remote host, and accept any certificate if the host sends one.
- *Optional trusted*: If the remote host sends a certificate, SSB checks if it is valid (not expired) and that the Common Name of the certificate contains the domain name or the IP address of the host. If these checks fail, SSB rejects the connection. However, SSB accepts the connection if the host does not send a certificate.



- **Optional untrusted:** Accept any certificate shown by the remote host. Note that the host must show a certificate.
- **Required trusted:** Verify the certificate of the remote host. Only valid certificates signed by a trusted certificate authority are accepted. See *Procedure 6.9.2, Uploading external certificates to SSB (p. 108)* for details on importing CA certificates. Note that the Common Name of the certificate must contain the domain name or the IP address of the host.
- **Required untrusted:** SSB requests a certificate from the remote host, and rejects the connection if no certificate is received; if the certificate is not valid (expired); or if the Common Name of the certificate does not contain the domain name or the IP address of the host.

**Note**

Consult the documentation of the remote server application to determine which protocols are supported.

UDP is a highly unreliable protocol and a high amount of messages may be lost without notice during the transfer. Use TCP or TLS instead whenever possible.

Step 5. Select the syslog protocol to use from the **Syslog protocol** field.

- To use the legacy BSD-syslog protocol described in RFC 3164, select **Legacy** and specify the message template to use. Select **Legacy** to use the message format described in the RFC; **ISO date** to replace the original timestamp with an ISO8061 compliant timestamp that includes year and timezone information. To customize the format of the message contents using macros, select **Custom message part only**, or **Custom on-wire message** to completely reformat the message (including the headers). For details on using macros, see *The syslog-ng Premium Edition Administrator Guide*. If you have no special requirements, use the **ISO date** template.
- Use the new IETF-syslog protocol. Note that most syslog applications and devices currently support only the legacy protocol. Consult the documentation of the remote server application to determine which protocols are supported. If you need, you can customize the contents of the message using macros. Note that for the IETF-syslog protocol, the header cannot be customized. For details on using macros, see *The syslog-ng Premium Edition Administrator Guide*.

Step 6. If SSB would send several messages with identical content to the destination, it can send only a single message and a line *Last message repeated n times..* Enter the number of seconds to wait for identical messages into the **Suppress timeout** field. This option corresponds to the *suppress()* parameter of syslog-ng.

Step 7. To limit the maximum number of messages sent to the destination per second, enter the maximum number of messages into the **Message throttle** field. Use this output-rate-limiting functionality only when using disk-buffer as well to avoid the risk of losing messages. Specifying 0 or a lower value sets the output limit to unlimited. This option corresponds to the *throttle()* parameter of syslog-ng.

Step 8. The timestamps of most log messages is accurate only to on second. SSB can include more accurate timestamps: set how many digits should be included in the **Timestamp fractions of a second** field. This option corresponds to the *frac_digits()* parameter of syslog-ng.



- Step 9. If the server and SSB are located in a different timezone and you use the *Legacy* message template (which does not include timezone information), select the timezone of the server from the **Timezone** field.
- Step 10. Set the size of the disk buffer in the **Output disk buffer** field. If the remote server becomes unavailable, SSB will buffer messages to the hard disk, and continue sending the messages when the remote server becomes available. This option corresponds to the *log_disk_fifo_size()* parameter of syslog-ng.
- Step 11. By default, SSB buffers up to 10000 messages in its memory if the remote server cannot accept them fast enough. To modify this value, adjust the **Output memory buffer** field as needed. This option corresponds to the *log_fifo_size()* parameter of syslog-ng.
- Step 12. Click **Commit**.
- Step 13. To start sending messages to the destination, include the new destination in a logpath. For details, see *Chapter 10, Managing log paths (p. 144)*.

9.4. Procedure – Forwarding log messages to SNMP destinations

Purpose:

To forward log messages from SSB to an SNMP destination, complete the following steps. The format of SSB SNMP messages conforms to the *CISCO-SYSLOG-MIB*.

Steps:

- Step 1. Navigate to **Log > Destinations** and select **+** to create a new remote destination.
- Step 2. Select **SNMP destination**.
- Step 3. Enter the IP address or hostname of the SNMP destination into the **Address** field. Enter the port where the server is accepting SNMP traps into the **Port** field.
- Step 4. Select the protocol version. The default value is *SNMP v2c*.
- To use the SNMP v2c protocol, select **SNMP v2c** and enter the name of the SNMP community to use in the **Community** field. The default value is *public*.
 - To use the SNMP v3 protocol, select **SNMP v3**. Enter the username and the Engine ID to be used when sending SNMP traps in the respective fields. Select the authentication method to use (MD5 or SH1) and enter the authentication password. Select the encryption method to use (Disabled or DES). In case of DES, enter the encryption password.
- Step 5. The timestamps of most log messages is accurate only to on second. SSB can include more accurate timestamps: set how many digits should be included in the **Timestamp fractions of a second** field. This option corresponds to the *frac_digits()* parameter of syslog-ng.
- Step 6. If the server and SSB are located in a different timezone and you use the *Legacy* message template (which does not include timezone information), select the timezone of the server from the **Timezone** field.
- Step 7. Set the size of the disk buffer in the **Output disk buffer** field. If the remote server becomes unavailable, SSB will buffer messages to the hard disk, and continue sending the messages when the remote server becomes available. This option corresponds to the *log_disk_fifo_size()* parameter of syslog-ng.



- Step 8. By default, SSB buffers up to 10000 messages in its memory if the remote server cannot accept them fast enough. To modify this value, adjust the **Output memory buffer** field as needed. This option corresponds to the `log_fifo_size()` parameter of syslog-ng.
- Step 9. Click **Commit**.
- Step 10. To start sending messages to the destination, include the new destination in a logpath. For details, see *Chapter 10, Managing log paths (p. 144)*.
- Step 11. To properly interpret and display the SNMP messages on your destination, download and install the [CISCO-SYSLOG-MIB](#) in your destination software.



Chapter 10. Managing log paths

This section describes how to create and configure log paths in SSB.

- For a list of default log paths, see *Section 10.1, Default logpaths in SSB (p. 144)*.
- For details on how to create a new log path, see *Procedure 10.2, Creating new log paths (p. 144)*.
- For details on how to send only selected messages to a destination, see *Section 10.3, Filtering messages (p. 146)*.

10.1. Default logpaths in SSB

Two log paths are available by default in SSB (see **Log > Paths**):

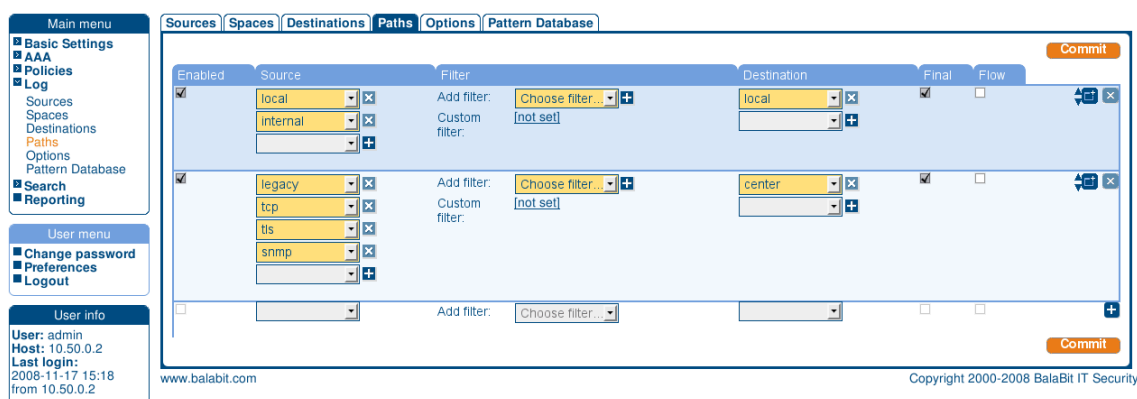


Figure 10.1. Default logpaths of SSB

- The first log path collects the local messages of SSB. It sends every message of the web interface, the built-in syslog-ng server, and other internal components to the **local** logspace.
- The second log path collects messages sent to SSB using the default syslog sources (for details, see *Section 7.1, Default message sources in SSB (p. 114)*) or via SNMP (for details, see *Procedure 7.2, Receiving SNMP messages (p. 114)*). These messages are stored in the **center** logspace.



Note

Note that both default log paths are marked as **Final**: if you create a new log path that collects logs from the default sources, make sure to adjust the order of the log paths, or disable the **Final** option for the default log path.

10.2. Procedure – Creating new log paths

Purpose:

To create a new log path, complete the following steps.

Steps:



- Step 1. Navigate to **Log > Paths** and select **+**. A new log path is added to the list of log paths.
- Step 2. Select a source for the log path from the **Source** field. Messages arriving to this source will be processed by this log path. To add more sources to the log path, select **+** in the source field and repeat this step.

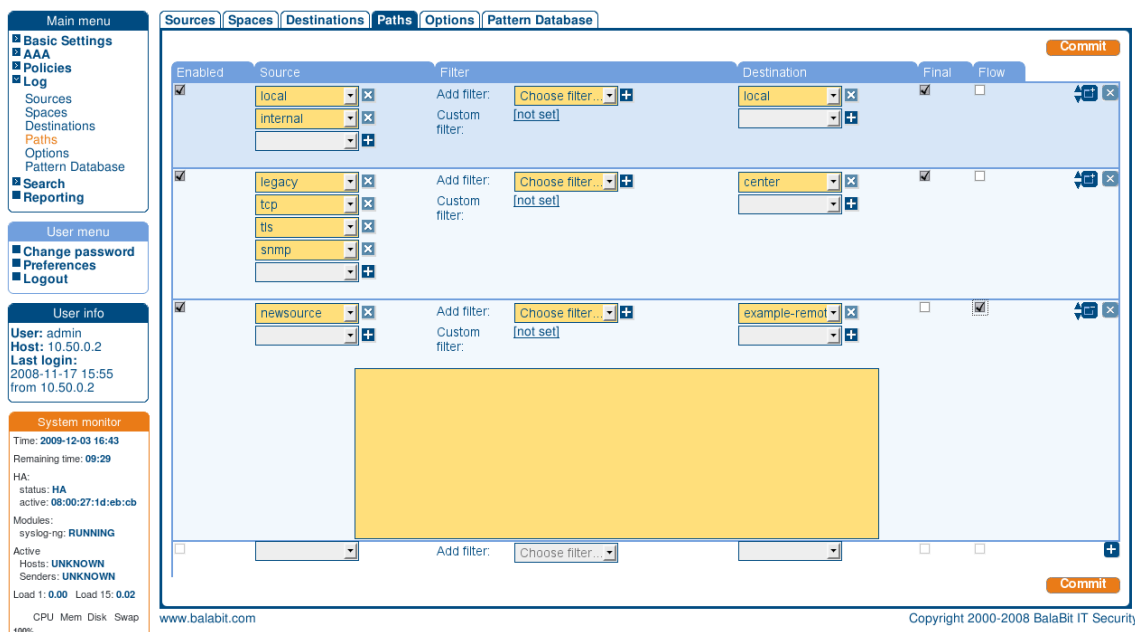


Figure 10.2. Creating a new logpath

Remote sources receive messages from the network, while *built-in* sources are messages that originate on SSB. However, note that the SNMP source (for details, see *Procedure 7.2, Receiving SNMP messages* (p. 114)) is listed in the built-in section.



Tip

To process every message of every source, leave the source option on *all*. This is equivalent to using the *catchall* flag of syslog-ng.

- Step 3. Select a destination for the log path from the **Destination** field. Messages arriving to this source will be forwarded to this destination. To add more destinations to the log path, select **+** in the destination field and repeat this step.



Note

Remote destinations forward the messages to external servers or databases and are configured on the **Log > Destinations** page (for details, see *Chapter 9, Forwarding messages from SSB* (p. 135)).

Local destinations store the messages locally on SSB and are configured on the **Log > Spaces** page (for details, see *Chapter 8, Storing messages on SSB* (p. 119)).

If you do not want to store the messages arriving to this log path, leave the **Destination** field on *none*.

**Warning**

The *none* destination discards messages — messages sent only to this destination will be lost irrevocably.

- Step 4. If you do not want other log paths to process the messages sent to a destination by this log path, select the **Final** option.
- The order of the log paths is important, especially if you use the **Final** option in one or more destinations, because SSB evaluates log paths in descending order. Use the buttons to position the log path if needed.
- Step 5. To enable flow-control for this log path, select the **Flow** option. For details on how flow-control works, see *Section 2.3, Managing incoming and outgoing messages with flow-control (p. 5)*.
- Step 6. If you do not want to send every message from the sources to the destinations, use filters. Select the filter to use from the **Filter** field, click , and configure the filter as needed. To apply more filters, click and select a new filter. Note that SSB sends only those messages to the destinations that pass every listed filter of the log path. The available filters are described in *Section 10.3, Filtering messages (p. 146)*.

Figure 10.3. Filtering log messages

- Step 7. Click **Commit**. After that, the new log path will start to collect log messages.

**Tip**

If you do not want to activate the log path immediately, deselect the **Enable** option.

10.3. Filtering messages

This section describes the filters that can be used in log paths. Every filter can be used to select (for example *priority is*) or exclude (for example *priority is not*) messages. The following filters are available:



- **facility:** Select messages sent by a specific facility (for example *kernel*).
- **host:** Select messages sent by a specific host. Enter the a hostname, IP address, or a POSIX (basic) regular expression.
- **message:** Select messages containing a specific keyword or POSIX (basic) regular expression in the text of the log message (excluding the headers).
- **priority:** Select messages of a specific priority.
- **program:** Select messages sent by a specific application. Enter the name of the application or a POSIX (basic) regular expression.
- **sender:** Filter on the address of the host that sent the message to SSB.



Note

The effect of the sender and the host filters is the same if every client sends the logs directly to SSB. But if SSB receives messages from relays, then the host filter applies for the address of the clients, while the sender applies for the address of the relays.

If multiple filters are set for a log path, only messages complying to every filter are sent to the destinations. (In other words, filters are added using the logical AND operation.)

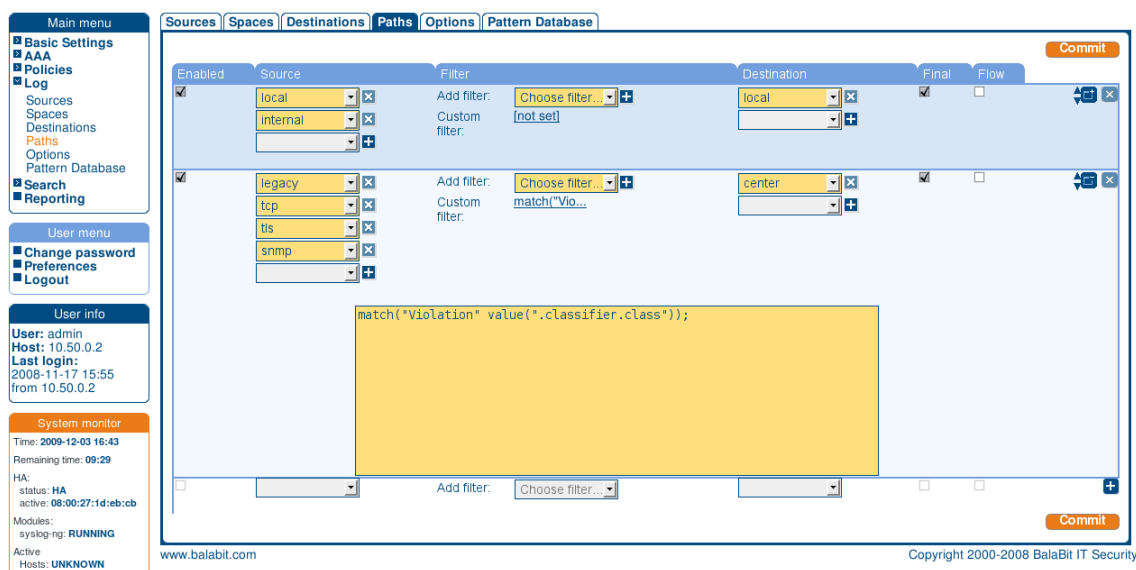


Figure 10.4. Using custom filters

If you need more complex filtering in your log path, select the of the log path and enter a custom filter into the appearing field. The contents of this field are pasted into the *filter()* parameter of the syslog-ng log path definition.

10.3.1. Procedure – Modifying messages using rewrite

Purpose:

The syslog-ng application can rewrite parts of the messages using rewrite rules. Almost all parts of the message can be rewritten. The rules use a key-value pair format.



Steps:

Step 1. Navigate to **Log > Paths**.

Step 2. Select the path(s) where you want to use rewrite rules.

Step 3. In the **Rewrites** section, click **+** to add a new rewrite rule. Rewrite rules can be applied before filtering, or after filtering.

The sequence of filtering and rewrite rules depends on how it was specified in the log path. The sequence of the process is the following:

1. Rewrite the message parts using the "before filtering" rewrite rules in the order the rewrite rules were given.
2. Filter the messages.
3. Rewrite the message parts using the "after filtering" rewrite rules in the order the rewrite rules were given.
4. Send the messages to the given destinations.

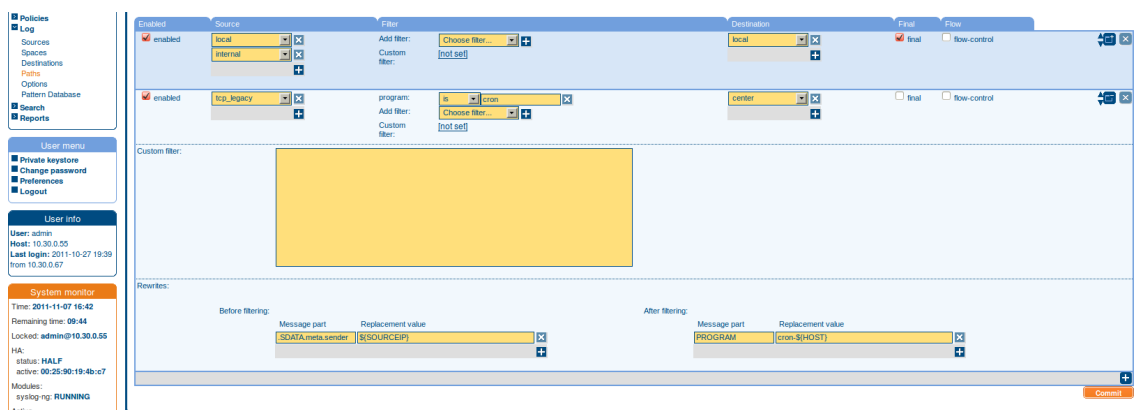


Figure 10.5. Modifying messages using rewrite

Step 4. Enter the part of the message to rewrite in the **Message part** field. For example `MESSAGE; HOST; .SDATA.meta.custom`. If the specified field does not exist, it is automatically created and set to the **Replacement value** field.

Step 5. Enter the value of the message part after rewriting in the **Replacement value** field. To use macros, begin with a \$ sign and enclose the name of the macro between braces, for example `${MSG}`; `${.SDATA.meta.custom}`.



Note

The replacement value completely replaces the old value of the message part.



Note

Hard macros contain data that is directly derived from the log message. It is not possible to change the values of hard macros in rewrite rules. For the list of hard macros, see *Section Hard vs. soft macros* in *The syslog-ng Premium Edition Administrator Guide*.



Chapter 11. Configuring syslog-ng options

There are several options of the syslog-ng server running on SSB that can be configured. These include:

- For details on general syslog-ng settings — see *Section 11.1, General syslog-ng settings (p. 149)*.
- For details on timestamping-related options — see *Section 11.2, Timestamping configuration on SSB (p. 151)*.
- For details on certificate management for receiving and sending log messages in TLS-encrypted channels — see *Procedure 11.4, Setting the certificates used in TLS-encrypted log transport (p. 153)*.
- For details on managing domain name resolution for log messages — see *Section 11.3, Using name resolution on SSB (p. 152)*.

11.1. General syslog-ng settings

To configure the general options of the syslog-ng server running on SSB, navigate to **Log > Options**. The following options are available (note that options related to name resolution are discussed in *Section 11.3, Using name resolution on SSB (p. 152)*):



The screenshot displays the 'Options' configuration page in the syslog-ng web interface. The left sidebar contains a 'Main menu' with links to Basic Settings, AAA, Policies, Log, Sources, Spaces, Destinations, Paths, Options, Pattern Database, Search, and Reports. Below this is a 'User menu' with links to Private keystore, Change password, Preferences, and Logout. The 'User info' section shows the user 'admin' with host '10.40.255.254' and last login '2011-11-02 13:43'. The 'System monitor' section shows the time '2011-11-02 14:28', remaining time '08:08', and system status 'syslog-ng: RUNNING'. The main configuration area is titled 'Options' and includes the following settings:

- Maximum logstore chunk time: 300 seconds
- Messages fetched in a single poll: 30 messages
- Initial window size: 100 messages
- Message size: 8192 bytes
- Wait time between polls: 0 milliseconds
- Idle time before destination is closed: 60 seconds
- DNS cache expiry: 3600 seconds
- Failed DNS cache expiry: 60 seconds
- Timestamp server: Local (selected) or Remote (Server URL: http://tsa.example.com)
- Timestamp policy OID: (empty)
- Cipher: aes-128-cbc
- Digest: sha1

Below the 'Options' section are several toggleable settings: TLS settings, SNMP source (checked), Alerting (checked), Artificial Ignorance (unchecked), Name resolving, Dashboard statistics (checked), and Message rate alerting statistics. A 'Commit' button is located at the bottom right of the configuration area.

Figure 11.1. Configuring syslog-ng options

- **Maximum logstore chunk time:** Time limit in seconds: syslog-ng closes the chunk if no new messages arrive until the time limit expires. Logstore chunks are closed when the time limit expires. If the time limit set in the **Idle time before destination is closed** option expires, the entire file is closed. This option corresponds to the `chunk_time()` parameter of syslog-ng.
- **Messages fetched in a single poll:** The maximum number of messages fetched from a source during a single poll loop. The destination queues might fill up before flow-control could stop reading if this parameter is too high. This option corresponds to the `log_fetch_limit()` parameter of syslog-ng.
- **Initial window size:** The size of the initial window used during flow control. This option corresponds to the `log_iw_size()` parameter of syslog-ng.
- **Message size:** Specifies the maximum length of incoming log messages in bytes. This option corresponds to the `log_msg_size()` parameter of syslog-ng.
- **Wait time between polls:** The time to wait in milliseconds before checking if new messages have arrived to a source. This option corresponds to the `time_sleep()` parameter of syslog-ng.
- **Idle time before destination is closed:** The time to wait in seconds before an idle destination file is closed. This option corresponds to the `time_reap()` parameter of syslog-ng.



11.2. Timestamping configuration on SSB

To configure the timestamping options of SSB, navigate to **Log > Options**. The following options are available:

- **Timestamp server:** Select the timestamping server to use for signing encrypted logspaces. To use the built-in timestamp server of SSB, select **Local**. To use an external timestamping server, select **Remote** and enter the address of the server into the **Server URL** field. Note that currently only plain HTTP services are supported, password-protected and HTTPS services are not supported at.



Warning

SSB currently supports only timestamping servers that use the [Internet X.509 Public Key Infrastructure Time-Stamp Protocol \(TSP\)](#) described in RFC 3161.

- **Timestamp policy OID:** If the Timestamping Server has timestamping policies configured, enter the OID of the policy to use into the Timestamping policy field. SSB will include this ID in the timestamping requests sent to the TSA.
- **Cipher:** Select the cipher method used to encrypt the logstore. The following cipher methods are available: *aes-128-cbc*, *aes-128-cfb*, *aes-128-cfb1*, *aes-128-cfb8*, *aes-128-ecb*, *aes-128-ofb*, *aes-192-cbc*, *aes-192-cfb*, *aes-192-cfb1*, *aes-192-cfb8*, *aes-192-ecb*, *aes-192-ofb*, *aes-256-cbc*, *aes-256-cfb*, *aes-256-cfb1*, *aes-256-cfb8*, *aes-256-ecb*, *aes-256-ofb*, *aes128*, *aes192*, *aes256*, *bf*, *bf-cbc*, *bf-cfb*, *bf-ecb*, *bf-ofb*, *blowfish*, *cast*, *cast-cbc*, *cast5-cbc*, *cast5-cfb*, *cast5-ecb*, *cast5-ofb*, *des*, *des-cbc*, *des-cfb*, *des-cfb1*, *des-cfb8*, *des-ecb*, *des-edc*, *des-edc-cbc*, *des-edc-cfb*, *des-edc-ofb*, *des-edc3*, *des-edc3-cbc*, *des-edc3-cfb*, *des-edc3-ofb*, *des-ofb*, *des3*, *desx*, *desx-cbc*, *rc2*, *rc2-40-cbc*, *rc2-64-cbc*, *rc2-cbc*, *rc2-cfb*, *rc2-ecb*, *rc2-ofb*, *rc4*, and *rc4-40*.

By default, SSB uses the *aes-128-cbc* method.

- **Digest:** Select the digest method to use. The following digest methods are available: *MD2*, *MD4*, *MD5*, *SHA-0* (*SHA*), *SHA-1*, *RIPEMD-160*, *SHA-224*, *SHA-256*, *SHA-384*, and *SHA-512*.

By default, SSB uses the *SHA-1* method.



Warning

The size of the digest hash must be equal to or larger than the key size of the cipher method. For example, to use the *aes-256-cbc* cipher method, the digest method must be at least *SHA-256*.

**Note**

The timestamp requests are handled by a separate process in syslog-ng; message processing is not affected if the timestamping server is slow or cannot be accessed.

11.3. Using name resolution on SSB

SSB can resolve the hostnames of the clients and include them in the log messages. However, the performance of SSB can be severely degraded if the domain name server is unaccessible or slow. Therefore, SSB automatically caches the results of name resolution. If you experience performance problems under high load, it is not recommended to disable name resolution. If you must use name resolution, consider the following:

The screenshot shows the SSB web interface. On the left, there's a sidebar with 'Main menu' (Basic Settings, AAA, Policies, Log, Sources, Spaces, Destinations, Paths, Options, Pattern Database, Search, Reports) and 'User menu' (Private key store, Change password, Preferences, Logout). Below that is 'User info' (User: admin, Host: 10.40.255.254, Last login: 2011-11-02 13:43 from 10.40.255.254) and 'System monitor' (Time: 2011-11-02 14:28, Remaining time: 08:08, Locked: admin@10.40.255.254, Modules: syslog-ng: RUNNING, Active Hosts: 2, Senders: 1, Load 1: 0.36 Load 15: 0.17, CPU Mem Disk Swap bar chart). The main area has tabs: Sources, Spaces, Destinations, Paths, Options, Pattern Database. The 'Options' tab is selected, showing a 'Commit' button at the top right. The 'Options' section includes: Maximum logstore chunk time: 300 seconds; Messages fetched in a single poll: 30 messages; Initial window size: 100 messages; Message size: 8192 bytes; Wait time between polls: 0 milliseconds; Idle time before destination is closed: 60 seconds; DNS cache expiry: 3600 seconds; Failed DNS cache expiry: 60 seconds; Timestamp server: Local (radio button), Remote (radio button, selected) with Server URL: http://tsa.example.com; Timestamp policy OID: (empty field); Cipher: aes-128-cbc; Digest: sha1. Below this is a 'TLS settings' section with a '+' icon. At the bottom, there's a list of checkboxes: SNMP source (checked), Alerting (checked), Artificial ignorance (unchecked), Name resolving (unchecked), Dashboard statistics (checked), and Message rate alerting statistics (unchecked). A 'Commit' button is at the bottom right.

Figure 11.2. Configuring DNS options

- If the IP addresses of the clients change only rarely, set the expiry of the DNS cache to a large value. By default, SSB caches successful DNS lookups for an hour, and failed lookups for one minute. These parameters can be adjusted under **Log > Options > Options > DNS Cache expiry** and **Failed DNS cache expiry**.
- Resolve the hostnames locally. Resolving hostnames locally enables you to display hostnames in the log files for frequently used hosts, without having to rely on a DNS server. The known IP address – hostname pairs are stored locally in a file. In the log messages, syslog-ng will replace the IP addresses of known hosts with their hostnames. To configure local name resolution, select **Log >**



Options > Name resolving, and enter the IP Address - hostname pairs in (for example `192.168.1.1 myhost.example.com`) into the **Persistent hostname list** field. Then navigate to **Log > Sources**, and set the **Use DNS** option of your sources to **Only from persistent configuration**.

The screenshot shows the SSB web interface. On the left is a sidebar with a 'Main menu' (Basic Settings, AAA, Policies, Log, Sources, Spaces, Destinations, Paths, Options, Pattern Database, Search, Reports) and a 'User menu' (Private keystore, Change password, Preferences, Logout). Below these is 'User info' (User: admin, Host: 10.40.255.254, Last login: 2011-11-02 14:22) and a 'System monitor' (Time: 2011-11-03 10:47, Remaining time: 07:50, Locked: admin@10.40.255.254, Modules: evl-in-m: RUNNING). The main content area has tabs for Sources, Spaces, Destinations, Paths, Options, and Pattern Database. The 'Options' tab is selected, showing a list of configuration items: Options, TLS settings, SNMP source (checked), Alerting (checked), Artificial Ignorance (checked), and Name resolving. The 'Name resolving' section is expanded, showing a 'Persistent hostname list' field with two entries: '192.168.1.1 myhost.example.com' and '192.168.1.2 myhost2.example.com'. At the bottom of the configuration area are 'Dashboard statistics' (checked) and 'Message rate alerting statistics', followed by a 'Commit' button.

Figure 11.3. Configuring persistent name resolution

11.4. Procedure – Setting the certificates used in TLS-encrypted log transport

Purpose:

To set a custom certificate and a CA certificate for encrypting the transfer of log messages, complete the following steps.



Note

If you do not upload a certificate to encrypt the TLS-communication (that is, the **TLS certificate** and **TLS private key** options are not set), SSB uses the certificate and CA certificate set for the web interface (set under **Basic Settings > Management > SSL certificates**) for this purpose as well.

Steps:

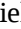


- Step 1. In your PKI system, generate and sign a certificate for SSB, then navigate to **Log > Options > TLS settings**.
- Step 2. Click the ☒ icon in the **TLS certificate** field to upload the certificate.



Figure 11.4. Configuring TLS settings for syslog-ng

To upload a certificate from a file, click **Browse** in the **Upload key** section, select the certificate file, and click **Upload**.

Alternatively, you can copy/paste the certificate into the **Key** field of the **Copy-paste key** section and click **Upload**.

- Step 3. Click the  icon in the **TLS private key** field and upload the private key corresponding to the certificate.
- Step 4. To set the certificate of the Certificate Authority (CA) used to verify the identity of the peers, click  in the **Certificate Authorities** field, then click .

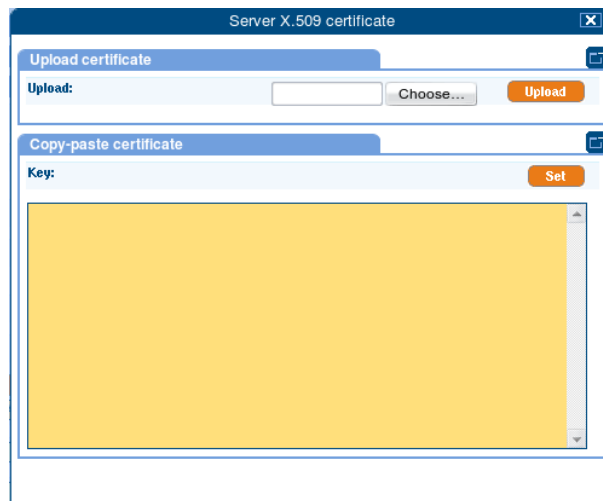


Figure 11.5. Uploading certificates

To upload a certificate from a file, click **Browse** in the **Upload key** section, select the certificate file, and click **Upload**.

Alternatively, you can copy/paste the certificate into the **Key** field of the **Copy-paste key** section and click **Upload**.

Repeat this step to add more CA certificates if needed.

- Step 5. If the CA issues a Certificate Revocation List (CRL), enter its URL into the **CRL URL** field. SSB periodically downloads the list and refuses certificates that appear on the list.



Note

Note that only *.pem* format CRLs are accepted. CRLs that are in PKCS7 format (*.crl*) are not accepted.

- Step 6. If you want to accept connections only from hosts using certain certificates signed by the CA, click **+** in the **Trusted distinguished names** field and enter the distinguished name (DN) of the accepted certificates into the **Distinguished name** field.
- Step 7. If you want to accept a certificate without uploading its corresponding CA certificate, click **+** in the **Trusted fingerprints** field and enter the SHA-1 fingerprint of the accepted certificates into the **SHA-1 fingerprint** field.



Chapter 12. Browsing log messages and SSB reports

This section describes how to browse the various types of log messages on SSB and exactly what kind of information do they contain.

- For the general use of the search interfaces, see *Section 12.1, Using the search interface (p. 156)*. Some of the search interfaces has certain special features, these are described in their respective sections.
- *Accounting information about the configuration changes performed on the SSB web interface*: Shows the activity of the SSB users and administrators. Available at **AAA > Accounting**. For the list of displayed parameters, see *Section 12.2, Changelogs of SSB (p. 159)*.
- *Collected log messages*: Log messages stored in a local log space or a remote database destination. To browse these messages, select **Search > Logs** and select the destination (that is, the logspace or database destination) from the **Destinations** field. For the list of displayed parameters, see *Section 12.3, Log messages collected on SSB (p. 160)*.
- *Peer configuration change*: Peers (client computers) that use syslog-ng Premium Edition 3.0 or newer send a special log message to SSB when their configuration is modified. These messages are located at **Search > Peer configuration change**. For the list of displayed parameters, see *Section 12.4, Configuration changes of syslog-ng peers (p. 175)*.
- *Alerts on the log messages*: If you use the pattern database of SSB to alert on certain log messages, then a history of the alerts is available at **Search > Alerts**. For the list of displayed parameters, see *Section 12.6, Log message alerts (p. 176)*.
- *Backup and archive notifications*: Notifications and errors encountered during backup or archiving are stored at **Search > Archive & Cleanup**. For the list of displayed parameters, see *Section 12.5, Notifications on archiving and backups (p. 176)*.
- *SSB reports*: PDF reports about the configuration changes, system health parameters, and other activities of SSB. Available at **Reporting > Reports**. For the list of displayed parameters, see *Section 12.8, Reports (p. 178)*.

12.1. Using the search interface

SSB has a uniform interface for browsing log messages, SSB configuration changes, reports, . This search interface consists of two main parts: a calendar bar and a table.

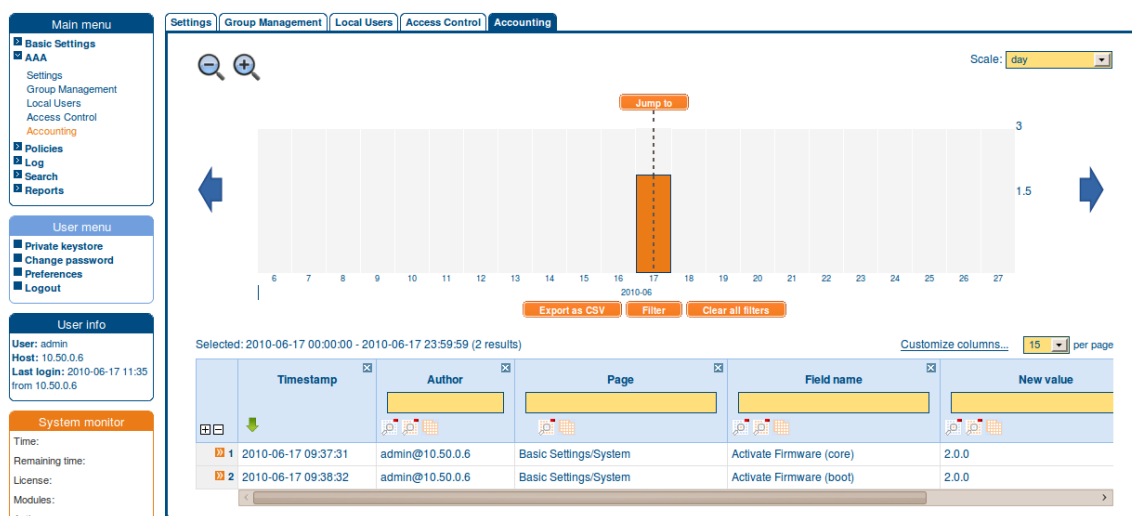


Figure 12.1. Using the search interface

The calendar bar displays the number of log messages in the selected interval. Use the and icons to zoom, and the arrows to display the previous or the next intervals. To explicitly select a date, select **Jump to** and set the date in the calendar. To select the length of the displayed interval use the **Scale** option.

Hovering the mouse above a calendar bar displays the number of entries and the start and end date of the period that the bar represents. Click a calendar bar to display the entries of that period in the table. Use Shift+Click to select multiple calendar bars. The **Selected** field shows the starting and ending date of the period listed in the table.

12.1.1. Procedure – Customizing columns

Purpose:

To select the data displayed on a search interface, complete the following steps:

Steps:

- Step 1. Navigate to the database you want to browse, for example **AAA > Accounting**.
- Step 2. Click **Customize Columns**. A popup window containing the list of visible and available columns and dynamic columns is displayed.

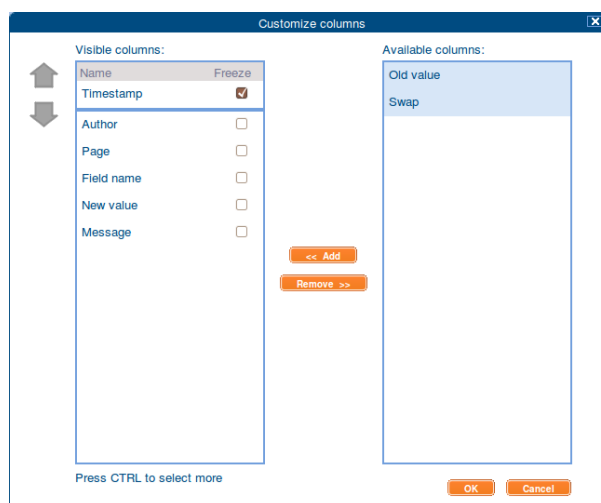


Figure 12.2. Displaying search information

Step 3. The displayed parameters are enlisted in the **Visible columns** field. All other available parameters are enlisted in the **Available columns** field.

- To add parameters to the **Visible columns** field, select the desired parameter(s) and click **Add**.
- To remove parameters from the **Visible columns** field, select the desired parameter(s) and click **Remove**.
- To freeze columns (to make them permanently visible, even when scrolling horizontally), enable the **Freeze** option next to the desired parameter.



Note

To select multiple parameters, press **Ctrl** while clicking the items.

Step 4. Click **OK**. The selected information is displayed.

12.1.2. Procedure – Adding and removing dynamic columns

Purpose:


Dynamic columns are created from name-value pairs. These columns are enlisted in the **Customize Columns** window (for details, see *Procedure 12.1.1, Customizing columns (p. 157)*), and can be searched with the help of the autocomplete function. To create (and remove) dynamic columns from name-value pairs, complete the following steps:

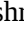

Steps:

Step 1. Name-value pairs are separated by commas in the **Details** column. Hover the mouse over a desired key in the name-value pair.

Step 2. Click .



- Step 3. The selected name generates a new, separate dynamic column with **<name>** heading (where **<name>** is the name of the key) before the **Details** column. The relevant values are displayed in the cells of the respective column.
- Step 4. To remove a dynamic column from the table, click  next to the heading. The key-value pair is moved back into the **Details** column.

If data is too long to fit on one line, it is automatically wrapped and only the first line is displayed. To expand a row, click . To shrink the row back to its original size, click . To expand/shrink all rows, click the respective button on the header of the table. The rows can also be expanded/shrunk by double clicking on the respective row.

The table can be filtered for any parameter, or a combination of parameters. To filter the list, enter the filter expression into the text box and press **Enter**, or click on an entry in the table. For example, to display only changes performed by a specific user, enter the username into the text box and press **Enter** — or just click on the specific username in the table.

**Note**

When you use filters, the calendar bar displays the statistics of the filtered results.

Filtering displays also partial matches: for example filtering **Author** for *adm* will display all changes performed by users whose username contains the *adm* string.

**Note**

Partial matching does not work in logstores; you can only search for complete tokens.

To save the table of search results as a file, click **Export as CSV**. This saves the table as a text file containing comma-separated values. Note that if an error occurs when exporting the data, the exported CSV file will include a line (usually as the last line of the file) starting with a zero and the details of the problem, for example *0;description_of_the_error*.

To restore the original table, click **Clear all filters**.

12.2. Changelogs of SSB

SSB automatically records the activity of its users and administrators. These activities are displayed at **AAA > Accounting**. The following information is available:

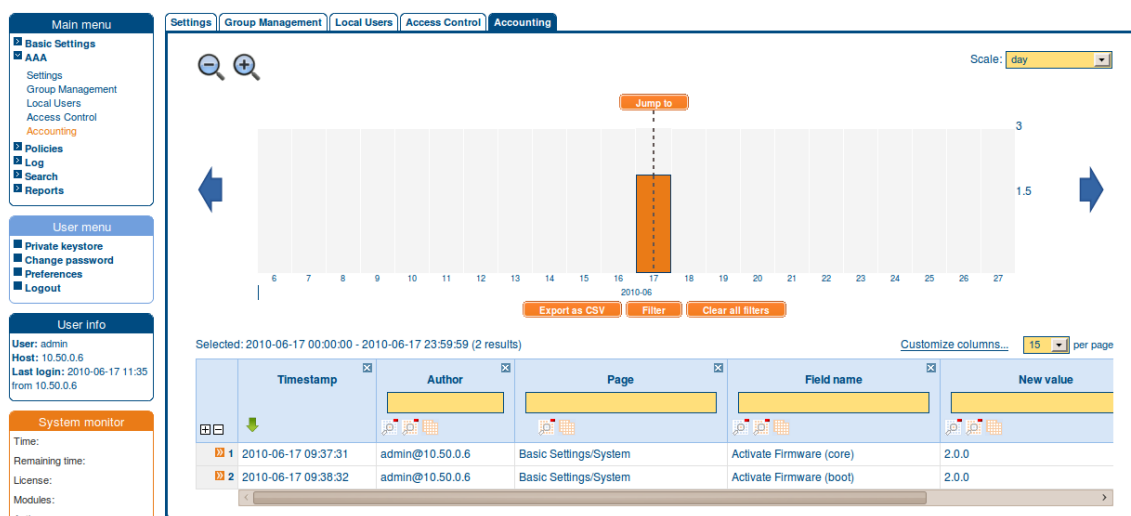


Figure 12.3. Displaying configuration changes

- **Timestamp:** The date when the modification was committed in *YEAR-MONTH-DAY HOUR:MINUTE:SECOND* format.
- **Author:** The SSB user who performed the modification.
- **Page:** The main menu item that was modified (for example *Basic Settings > Management*).
- **Field name:** The name of the field on the page that was modified.
- **New value:** The new value of the field after the modification.
- **Description:** The changelog entered by the SSB administrator. Changelogs are available only if the **AAA > Settings > Require commit log** option was enabled at the time of the change.
- **Old value:** The original value of the field.
- **Swap:** Signs if the order of objects was modified on the page (for example the order of two policies in the list).

12.3. Log messages collected on SSB

Log messages stored locally on SSB or in a remote database can be browsed at **Search > Logs**. Select the destination storing the logs you want to browse from the **Destination** field.

- For a description about the available columns, see *Section 12.3.1, Metadata collected about log messages (p. 161)*.
- For details about how to use and save filters, see *Section 12.3.3, Using and managing search filters (p. 166)*.
- For details about how to display statistics about your search results, see *Procedure 12.3.4, Displaying statistics on search results (p. 168)*.
- For details about how to browse and search log messages that are stored in an encrypted logstore, see *Section 12.3.5, Browsing encrypted log spaces (p. 171)*.



12.3.1. Metadata collected about log messages

The following information is available about the log messages:

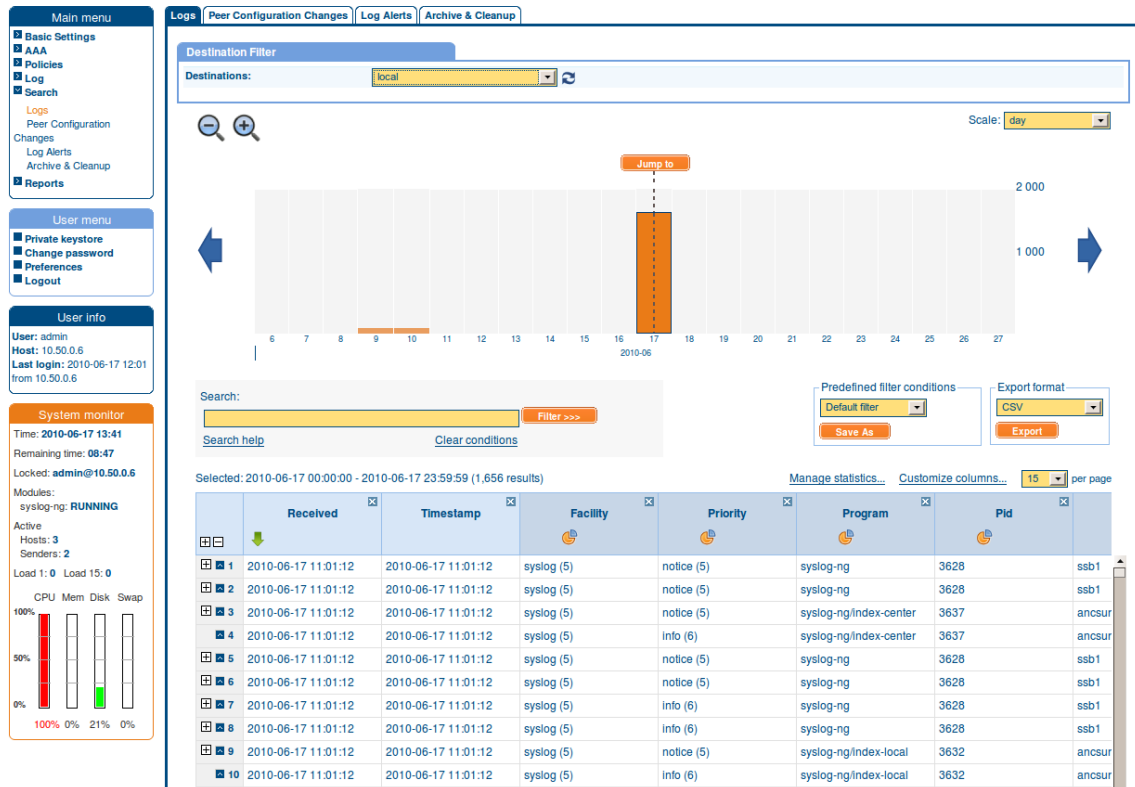


Figure 12.4. Displaying search information

The columns that are indexed and can be searched are marked with a darker blue table heading.

- **Received:** The date when SSB has received the log message in *YEAR-MONTH-DAY HOUR:MINUTE:SECOND* format.
- **Timestamp:** The timestamp received in the message — the time when the log message was created in *YEAR-MONTH-DAY HOUR:MINUTE:SECOND* format.
- **Facility:** The facility that sent the message.
- **Priority:** The priority value of the message.
- **Program:** The application that created the message.
- **Pid:** The program identifier of the application that created the message.
- **Host:** The IP address or hostname of the client that sent the message to SSB.
- **Message:** The text of the log message.
- **Tag:** Tags assigned to the message matching certain pattern database rules.
- **Rule ID:** ID of the pattern database rule that matched the message.
- **Details:** Name-value pairs assigned to the message.

**Note**

Note that SSB only indexes the first 60 characters of every name of word. This means that if the name of a name-value pair is longer than 60 characters, searching for the exact name will not give any results. For example, if the name of a name-value pair is `.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-12345`, SSB will index only the first 60 characters: `.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-`. Searching for `.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-12345`, will not return any results, so search for `.sdata.security.uid=2011-12-08T12:32:25.024+01:00-hostname-` (or a part of it) instead.

- *Rule description*: Description of the pattern database rule that matched the message.

12.3.2. Using wildcards and boolean search

Searching is facilitated in SSB by the ability to search for wildcards and boolean expressions. The following sections provide examples for different search queries.

- For details about exact match and complex queries, see *Section Searching exact match and complex queries (p. 162)*
- For details about searching in a specific part of the message, see *Section Searching in a specific part of the message (p. 163)*
- For details about combining search keywords, see *Section Combining search keywords: (p. 163)*
- For details about using wildcard searches, see *Section Using wildcard searches: (p. 163)*
- For details about searching for special characters, see *Section Searching for special characters : (p. 165)*

Searching exact match and complex queries

By default, SSB searches for keywords as whole words in the MESSAGE part of the log message and returns only exact matches.

**Example 12.1. Searching exact match**

For example, searching for the *example* keyword:

Search expression:

- example

Matches:

- example

Does not match:

- examples
- example.com
- query-by-example
- exam

The specific part of the message to perform the search on can be given, more complex search queries can be written using boolean expressions and wildcard search is supported in all the fields:

**Example 12.2. Complex search queries**

Search expression:

- message1 AND (host:testhost OR program:syslog*)

Searching in a specific part of the message

You can search in a specific part of the message using the `<type>:` prefix. The `message:` (or `msg:`) prefix means the message part and can be omitted. For example, to search for the name of an application, use the `program:` prefix. To search for a host name, use the `host:` prefix, and so on.

Combining search keywords:

You can use boolean operators to combine search keywords. More complex search expressions are also available, by using parentheses.

**Example 12.3. Combining keywords in search**

You can combine keywords with the AND, OR, NOT operators, for example:

- keyword1 AND keyword2 (returns log messages that contain both keywords)
- keyword1 OR keyword2 (returns log messages that contain at least one of the keywords)
- host:host1 AND program:prg1
- message:example AND program:prg1

To search for expressions that can be interpreted as boolean operators (for example AND), use the following format: `message:AND`.

**Example 12.4. Using parentheses in search**

Use parentheses to create more complex search expressions, for example:

- (keyword1 OR keyword2) AND keyword3

Using wildcard searches:

You can use the following wildcard characters in your search expressions:

**Example 12.5. Using wildcard ? in search**

? (question mark) means exactly one arbitrary character, for example:

- Search expression:
 - example?
 Matches:
 - example1
 - examples
 Does not match:
 - example.com
 - example12
 - query-by-example
 - example?
- Search expression:
 - ?example?



Matches:

- 1example2

Does not match:

- example.com
- example12
- query-by-example

■ Search expression:

- example??

Matches:

- example12

Does not match:

- example.com
- example1
- query-by-example

The ? wildcard search does not work for finding non-UTF-8 or multibyte characters. If you want to search for these characters, the expression ?? might work. The * wildcard search finds non-UTF-8 and multibyte characters as well.



Example 12.6. Using wildcard * in search

* means 0 or more arbitrary characters, for example:

■ Search expression:

- example*

Matches:

- example
- examples
- example.com

Does not match:

- query-by-example
- example*

■ Search expression:

- *example

Matches:

- example
- query-by-example

Does not match:

- example.com
- example12

■ Search expression:

- *example*

Matches:

- example
- query-by-example
- example.com
- example12

The * wildcard search finds non-UTF-8 and multibyte characters as well.

**Example 12.7. Using combined wildcards in search**

Wildcard characters can be combined, for example:

Search expression:

- `ex?mple*`

Matches:

- `example1`
- `examples`
- `example.com`
- `exemple.com`

Does not match:

- `exmples`
- `example12`
- `query-by-example`

Wildcard characters also work in any message part, for example, `program:postfix*`.

Searching for special characters :

To search for the question mark (?), asterisk (*), backslash (\) or whitespace () characters, you must prefix these characters with a backslash (\). Any character after a backslash is handled as character to be searched for. For example:

**Example 12.8. Searching for special characters**

To search for a special character, use a backslash (\).

Search expression:

- `example\?`

Matches:

- `example?`

Does not match:

- `examples`
- `example1`

To search for a special character, backslash character, use two backslashes (\\).

Search expression:

- `C:\\Windows`

Matches:

- `C:\Windows`

Search expression:

- `nvpair:path=C:\\Program\ Files`

Matches:

- `C:\Program Files`

**Note**

It is not possible to search for whitespace () character in the MESSAGE part of the log message, since it is a hard-coded delimiter character.

12.3.3. Using and managing search filters

- To filter the search results, set the filters you need and click **Filter**.

When typing search expressions, you can create complex searches by using the AND, OR, and NOT operations. Note that whitespace is interpreted as AND, and NOT can be used only together with other operators (for example *adm AND NOT admin*). Also, you can search in the various fields using the *fieldname:keyword* method, for example to search the **Program** field for entries of the *CRON* application, use *program:CRON*.

When searching log messages, the capabilities of the search engine depend on the delimiters used to index the particular log space. For details on how to configure the delimiters used for indexing, see *Procedure 8.4.1, Creating a new logstore (p. 121)*.

- To apply a predefined filter, select the filter from the **Predefined filter conditions** field.
- To delete a predefined filter, select the filter from the **Predefined filter conditions** field and click **Delete**. Note that you need the **Manage global filters** privilege to delete global filters.
- To create and save a filter, complete *Procedure 12.3.3.1, Creating and saving filters for later use (p. 166)*.

12.3.3.1. Procedure – Creating and saving filters for later use

Purpose:

To create and save a filter for later use, complete the following steps:

Steps:

- Step 1. Navigate to **Search > Logs**, and select a destination.
- Step 2. Set the filters you need.
- Step 3. Select **Predefined filter conditions > Save As**. A popup window is displayed.
- Step 4. Enter a name for the filter into the **Name** field.



Figure 12.5. Saving filter conditions

Step 5. If you want the filter to be available for other SSB users as well, select **Global**. To restrict the availability of the filter to a set of specific users, use the **Group** field. **Local** filters are visible only for you.



Note

Adding a search filter to a group or user with no Search rights grants limited access to the search interface: they can view (and search within) the results of the assigned search filters only.

Users or groups with Search rights are able to see all global Search filters, regardless of the assigned usergroup(s).



Note

Filters cannot be modified later, only deleted. A filter can be deleted by the user who created it, and by users whose group has the **Search > Manage global filters** privilege.

Step 6. To modify the timeframe of the search, select **Interval**, and set the beginning and ending date and time of the search. This is useful when you want to display only the logs of a specific event. Note that you must always set an interval for global filters.

Step 7. Click **OK**.



Note

When searching log messages, the capabilities of the search engine depend on the delimiters used to index the particular log space. Starting with SSB 1.0.2, it is possible to configure the delimiters used for indexing, but it is not possible to re-index messages already received. This means that prior to version 1.0.2, IP addresses, MAC addresses, and similar information cannot be automatically retrieved from the text of the log message, because the . and : characters were used as delimiters, and the segments of the IP and MAC addresses are treated as separate tokens. For details on how to configure the delimiters used for indexing, see *Procedure 8.4.1, Creating a new logstore (p. 121)*.



12.3.4. Procedure – Displaying statistics on search results

Purpose:

SSB can create statistics (bar, pie and list) from various information about the search results, for example, the distribution of the sender hosts, and so on. To display statistics about the log messages, complete the following steps:

Steps:

- Step 1. Navigate to **Search > Logs**, and select a destination.
- Step 2. Set the filters you need.
- Step 3. Click a pie chart icon in the header of the table. A popup window is displayed.

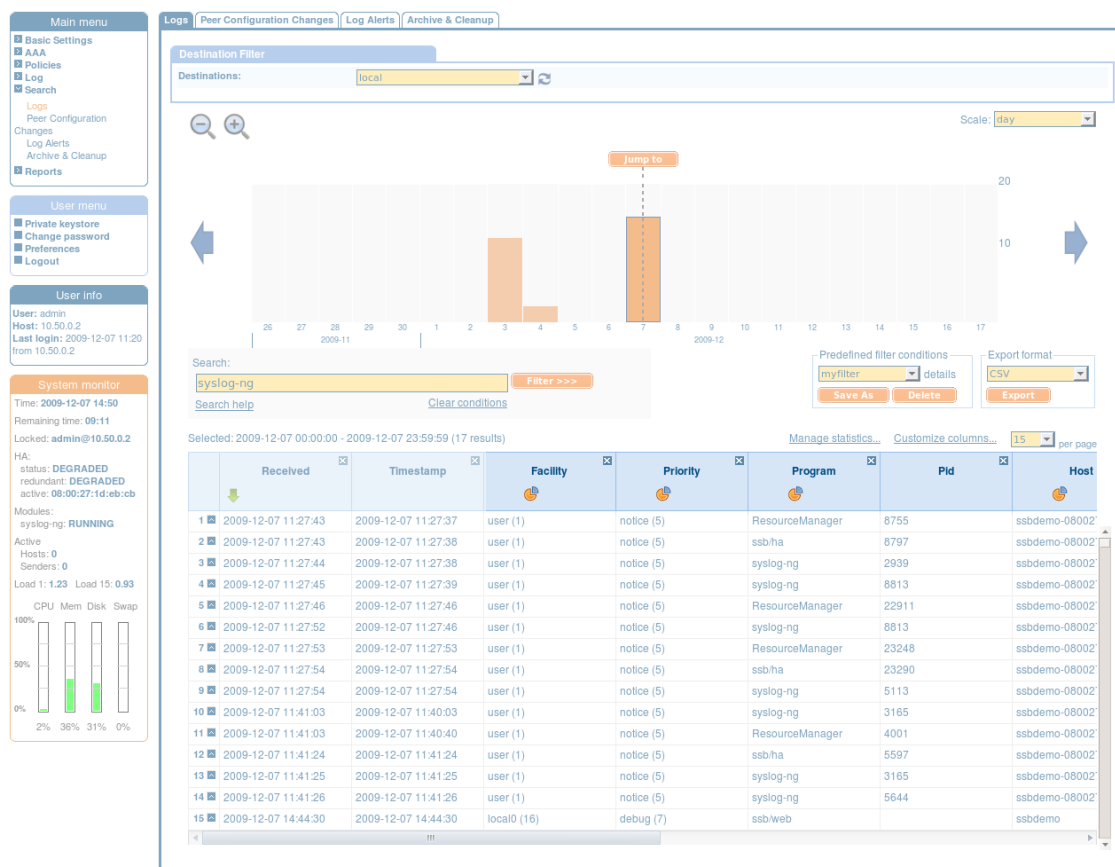


Figure 12.6. Displaying statistics

- Step 4. Select the type of metadata you want to create statistics on from the **Statistics based on** field. Currently you can create statistics from the *Facility*, *Priority*, *Program*, *Pid*, *Host*, *Tags*, and *.classifier.class* fields.

**Note**

The `.classifier.class` data is the class assigned to the message when pattern database is used. For details, see *Chapter 13, Classifying messages with pattern databases (p. 183)* The pattern databases provided by BalaBit currently use the following message classes by default: *system*, *security*, *violation*, or *unknown*.

Step 5. Select the type of chart to display, that is, **Bar**, **Pie** or **List**. The chart will be displayed in the same popup window.

**Note**

In **List** view, percentages add up to 100%. The only exception to this is when statistics are based on **Tags**. Since it provides statistics for tags rather than messages, it is possible that if messages have multiple tags, the percentages will add up to more than 100%.

Step 6. By default, the statistics start with the largest number of entries. To start statistics with the least number of entries, select **Least**.

Step 7. Select the number of data groups to display from the **Number of entries** field. For example, if you want to display the statistics of the ten hosts that send the most messages (the "top talkers"), select **10**. That way the top ten talkers will be displayed individually, while the amount of messages sent by the other hosts will be aggregated and labeled as **Others**.

**Note**

For pie and bar charts you can select 5, 10 and 15, for lists 5, 10, 15, 50 and 100.




Figure 12.7. Selecting display type

Step 8. *Optional step:* To export statistics data to a CSV file, select **List**, set the number of entries and click **Export all to CSV**. SSB compiles the selected data into a `results.csv` file.



Note

This action exports all rows, not only the currently displayed ones.

Step 9. *Optional step:* You can also save these statistics and include them in reports as a report subchapter. You can include these subchapters into your reports in the **Reports > Configuration** menu.

Step a. To save these statistics as custom statistics for reporting, click **Save As Custom Statistics For Reporting**.

Step b. Add a name for the statistics in the **Name** field.

Step c. Select a group from the already existing groups in the **Group** field. The autocomplete function helps you with the selection.

Step d. *Optional step:* The **Add to report as a subchapter** function enables you to instantly add this statistics as a subchapter to the selected report.



Step e. Click **Save**. This action includes the saved statistics as a selectable subchapter into **Reports > Configuration**. For details on how to add this subchapter to a selected report, see *Procedure 12.8.2, Configuring custom reports (p. 180)*.

12.3.5. Browsing encrypted log spaces

By default, you cannot browse encrypted logstores from the SSB web interface, because the required decryption keys are not available on SSB. To make browsing and searching encrypted logstores possible, SSB provides the following options:

- *Upload the required decryption keys and make them available to every SSB user:* It is possible to upload the private key (or set of keys) to SSB, and use these keys to decrypt the logstore files. That way anyone who has the right to search the particular log space can search the messages, but that also means that the decryption keys are directly stored unencrypted in the SSB configuration file. As this may raise security concerns, avoid this solution unless absolutely necessary. For details, see *Procedure 12.3.5.1, Assigning decryption keys to a logstore (p. 171)*.
- *Upload the decryption keys for a single user:* It is possible to upload decryption keys and bind them to a user account. That way, only this particular user can search the logstores encrypted with the certificate corresponding to the private key. The decryption keys will be stored on SSB but they are only available for the particular user, and can be encrypted (requiring a passphrase to access). For details, see *Procedure 12.3.5.2, Assigning decryption keys to a user account (p. 172)*.
- *Upload keys temporarily:* Users can upload decryption keys to SSB temporarily: they upload the keys, browse the log messages, and the keys are automatically deleted when the user logs out from SSB. For details, see *Procedure 12.3.5.3, Using temporal decryption keys (p. 174)*.

**Note**

Do not use SSB's own keys and certificates for encrypting or decrypting.

12.3.5.1. Procedure – Assigning decryption keys to a logstore

Steps:


- Step 1. To upload the decryption keys, navigate to **Log > Spaces** and select the encrypted logspace you want to make searchable for every user via the SSB web interface.
- Step 2. Select **Decryption private keys** > . A popup window is displayed.



Figure 12.8. Adding decryption keys to a logstore

- Step 3. Paste or upload the private key of the certificate used to encrypt the logstore.
- Step 4. Repeat Steps 2-3 to upload additional keys if needed.
- Step 5. Click **Commit**.

12.3.5.2. Procedure – Assigning decryption keys to a user account

Steps:



- Step 1. To upload the decryption keys to a specific user account, login to SSB, and select **User Menu > Private keystore**. A popup window is displayed.
- Step 2. Select **Permanent** > , then select **Certificate** > . A popup window is displayed.



Figure 12.9. Adding decryption keys to the private keystore

- Step 3. Paste or upload the certificate used to encrypt the logstore.
- Step 4. Select **Key** > ☒. A popup window is displayed.
- Step 5. Paste or upload the private key of the certificate used to encrypt the logstore.
- Step 6. Repeat Steps 2-5 to upload additional keys if needed.
- Step 7. Select **Security passphrase** > **Change**, and enter a passphrase to protect the private keys.

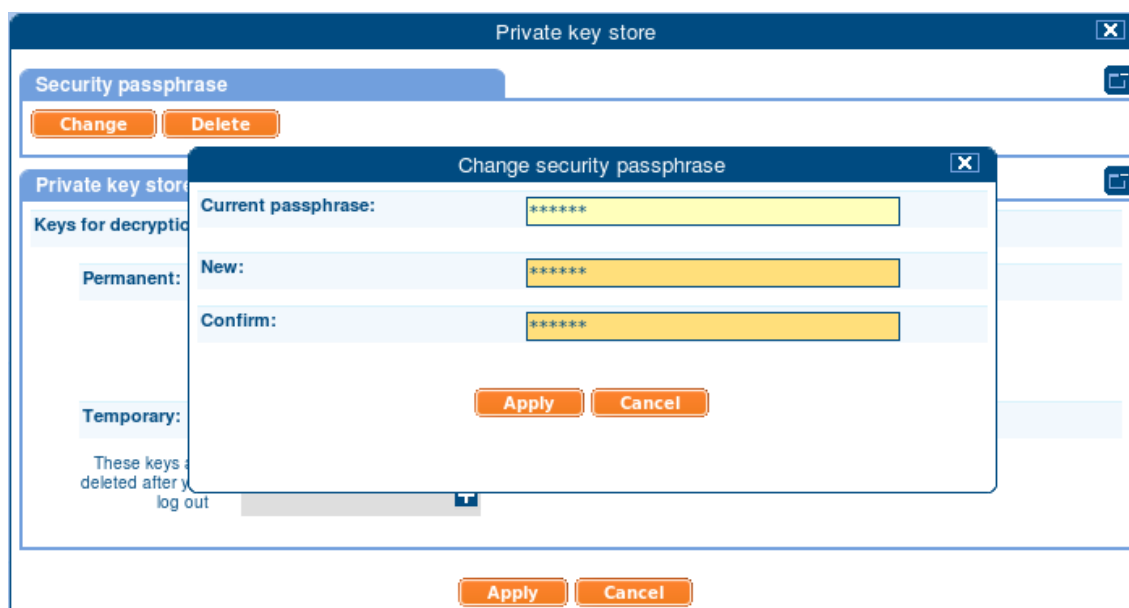


Figure 12.10. Adding decryption keys to the private keystore

Step 8. Click **Apply**.

12.3.5.3. Procedure – Using temporal decryption keys

Steps:

- Step 1. To upload the decryption keys temporarily, login to SSB, and select **User Menu > Private keystore**. A popup window is displayed.
- Step 2. Select **Temporary > +**, then select **Certificate > [icon]**. A popup window is displayed.

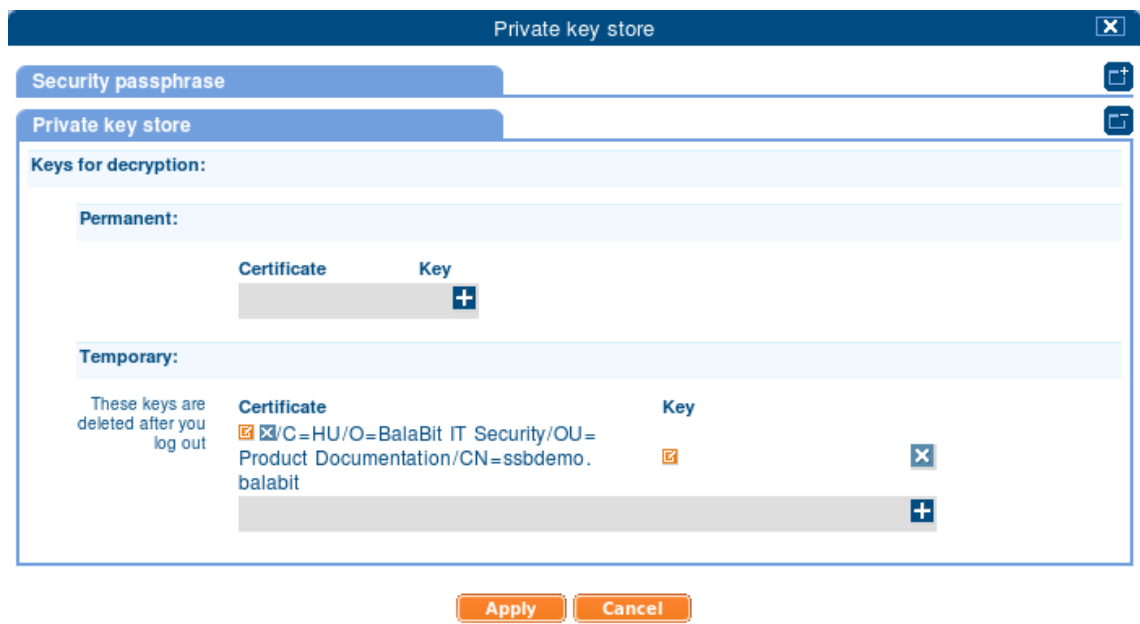


Figure 12.11. Adding decryption keys to the private keystore

- Step 3. Paste or upload the certificate used to encrypt the logstore.
- Step 4. Select **Key** > ☒. A popup window is displayed.
- Step 5. Paste or upload the private key of the certificate used to encrypt the logstore.
- Step 6. Repeat Steps 2-5 to upload additional keys if needed.
- Step 7. Click **Apply**.

12.4. Configuration changes of syslog-ng peers

Peers running syslog-ng Premium Edition 3.0 or later automatically send a notification to SSB when their configuration has changed since the last configuration reload or restart. These log messages are available at **Search > Peer Configuration Change**. Note that the log messages do not contain the actual modification; only indicate that the configuration was modified. The following information is available:

- **Timestamp:** The timestamp received in the message — the time when the log message was created in `YEAR-MONTH-DAY HOUR:MINUTE:SECOND` format.
- **Hostname:** The hostname or IP address of the client whose configuration has been changed.
- **Validity:** The validation of the checksum signature.
- **Version:** Version number of the syslog-ng application that sent the message.
- **Sender address:** The IP address of the client or relay that sent the message directly to SSB.
- **Signature:** The signature of the syslog-ng client.
- **Fingerprint:** The SHA-1 hash of the new configuration file.



12.5. Notifications on archiving and backups

Notifications and error messages of the archiving, cleanup and backup procedures are available at **Search > Archive & Cleanup**. The following information is available:

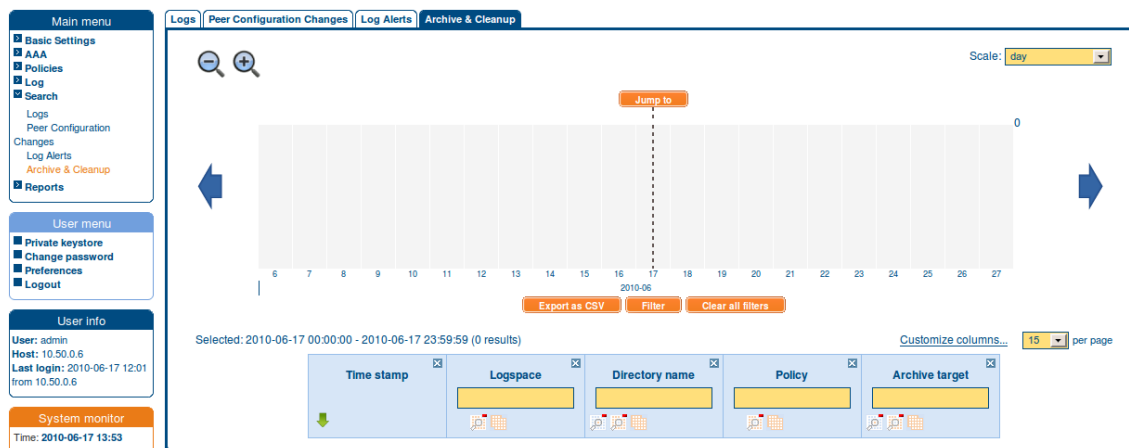


Figure 12.12. Displaying archiving and backup notifications

- **Timestamp:** The date of the message in in *YEAR-MONTH-DAY HOUR:MINUTE:SECOND* format.
- **Logspace:** Name of the archived of backed up logspace.
- **Directory name:** .
- **Policy:** Name of the archive or backup policy used.
- **Archive target:** Address of the remote server used in the policy.
- **Manual archiving:** Indicates if the archiving or backup process was started manually.

12.6. Log message alerts

When using the pattern database, SSB raises alerts for messages that are classified as *Violation*. The history of these alerts is available at **Search > Alerts**. The following information is available about the alerts:

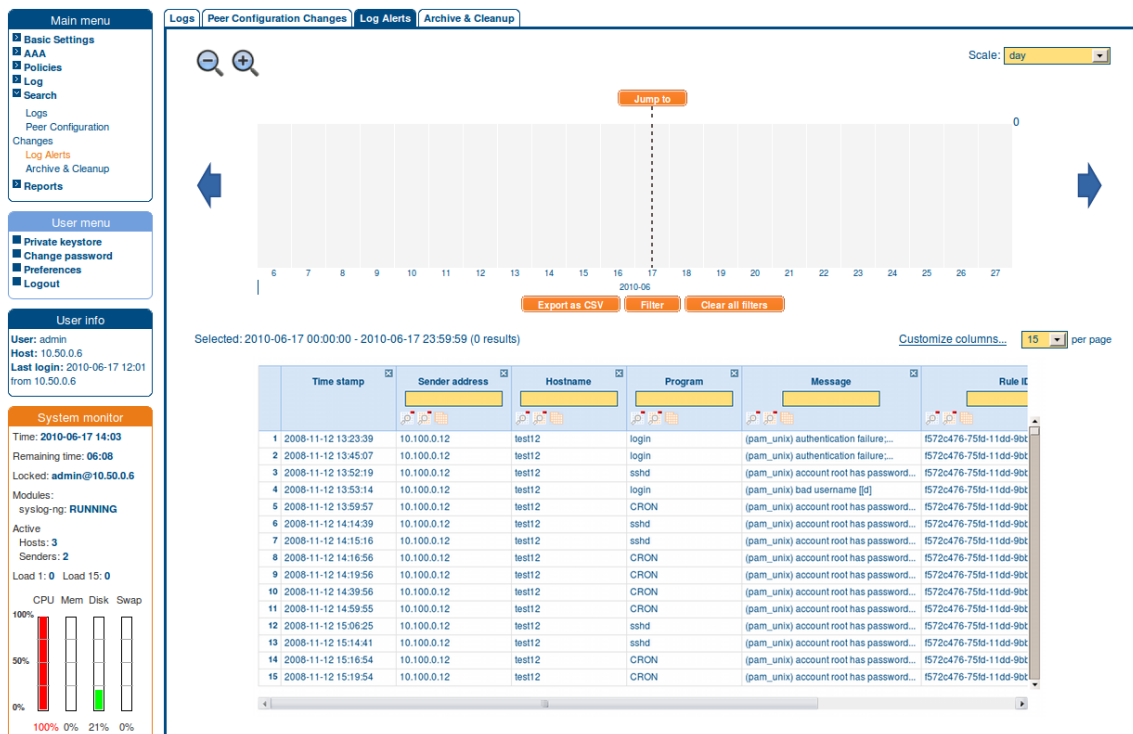


Figure 12.13. Displaying alert messages

- **Timestamp:** The date of the alert in `YEAR-MONTH-DAY HOUR:MINUTE:SECOND` format.
- **Sender address:** The IP address of the client or relay that sent the message directly to SSB.
- **Hostname:** The hostname or IP address of the client that sent the message.
- **Program:** The application that generated the message.
- **Message:** The content of the message.
- **Rule ID:** The ID of the classification rule in the pattern database that matched the message. For details, see *Chapter 13, Classifying messages with pattern databases (p. 183)*.
- **Rule description:** The description of the classification rule that matched the message. For details, see *Chapter 13, Classifying messages with pattern databases (p. 183)*.

12.7. Statistics collection options

To control the quantity and quality of the statistics collected to the **Dashboard** (for details, see *Section 14.5, Status history and statistics (p. 196)*), set the statistics collection options.

Navigate to **Log > Options > Dashboard Statistics**.

Time-based statistics: the default setting is **Enabled**.

- **Cleanup if unchanged for:** Statistics unchanged (not present in syslog-ng statistics output any more) for this number of days will be cleaned up from the system. Enter `0` here to keep them forever. To start the cleanup process immediately, click **Cleanup now**.



- **Enable statistics for:** the default setting is that all checkboxes are enabled. This allows you to select which options to collect statistics for. To display the collected statistics for an option, navigate to **Basic Settings > Dashboard > Syslog-ng statistics**, select **Time-based statistics** and select the desired option.

**Note**

When disabling an option, the data will only be deleted after the first cleanup. Until then, the already collected data is still accessible on the dashboard.

Top/Least statistics: the default setting is **Enabled** and all checkboxes are enabled. This allows you to select which options to collect statistics for. To display the collected statistics for an option, navigate to **Basic Settings > Dashboard > Syslog-ng statistics**, select **Top/Least statistics** and select the desired option.

Maximum number of statistics to process: Enter the number of statistics files to keep on the system. Enter 0 here to store unlimited number of statistics files. Statistics over this limit will be dropped, and SSB sends an error message containing the number of entries dropped and the first dropped entry. This setting needs to be increased only if you have more than 10000 hosts.

Sampling interval: Select the sampling interval for the statistics here. A more frequent sampling interval results in more precise graphs at the cost of heavier system load. The default setting is *5 minutes*. The possible parameters are *5 minutes, 10 minutes, 30 minutes, 60 minutes, 2 hours, 4 hours, 8 hours, 1 day*.

**Warning**

Hazard of data loss! When changing the Sampling interval, the already existing statistics are not converted to the new sampling rate, but are deleted.

To clear all statistics, click **Clear all statistics**. It is advised to clear statistics if you have changed the number of the statistics files to keep, or if you have disabled the time-based statistics collection.

12.8. Reports

SSB periodically creates reports on the activity of the administrators, the system-health information of SSB, as well as the processed traffic. These reports are available in Portable Document (PDF) format by selecting **Reports > Generated reports** from the Main Menu. The reports are also sent to the e-mail address set at **Basic Settings > Management > Mail settings > Send reports to**, unless specified otherwise in the configuration of the report.

To access the reports from the SSB web interface, the user must have the appropriate privileges.

**Note**

If the **Basic Settings > Management > Mail settings > Send reports to** address is not set, the report is sent to the SSB administrator's e-mail address.

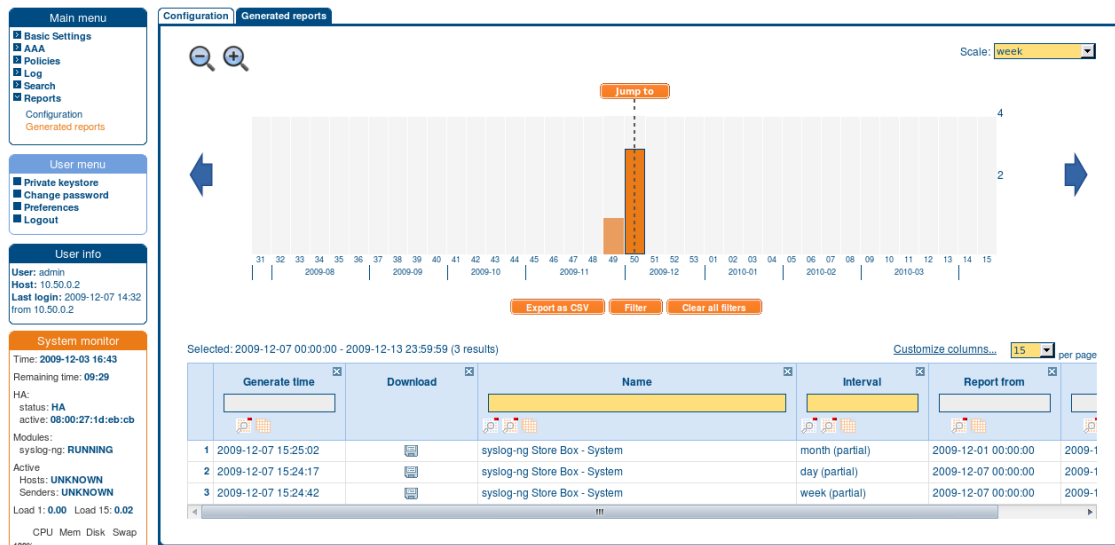


Figure 12.14. Browsing reports

Reports are generated as follows:

- **Daily reports** are generated every day at 00:01.
- **Weekly reports** are generated every week on Monday at 00:01.
- **Monthly reports** are generated on the first day of every month at 00:01.



Tip

Use the time bar to find reports that contain a particular period. If you select a period (for example click on a bar), only those reports will be displayed that contain information about the selected period.

The following information is available about the reports:

- **Download:** A link to download the report.
- **Name:** Name of the report.
- **Interval:** The length of the reported period, for example week, month, and so on.
- **Report from:** The start of the reported interval.
- **Report to:** The end of the reported interval.
- **Generate time:** The date when the report was created.



Tip

To create a report for the current day, select **Generate reports for today**. The report will contain data for the *00:00 - current time* interval. If artificial ignorance (for details, see Chapter 13, *Classifying messages with pattern databases* (p. 183)) is enabled, an artificial ignorance report is created as well.



12.8.1. Contents of the default reports

The default report of SSB (called *System*) is available in Adobe Portable Document Format (PDF), and contains the following information for the given period:

- **Configuration changes:** Lists the number of SSB configuration changes per page and per user. The frequency of the configuration changes is also displayed on a chart.
- **Peer configuration:** Lists the number of times the configuration of a syslog-ng client was changed per client, as well as the version number of the syslog-ng application running on the client (if this information is available).
- **Alerts:** Various statistics about the alerts received from classifying messages using the pattern database (if pattern databases have been uploaded to SSB).
- **syslog-ng traffic statistics:** Displays the rate of incoming, forwarded, stored, and dropped messages in messages/second.
- **System health information:** Displays information about the filesystem and network use of SSB, as well as the average load.

12.8.2. Procedure – Configuring custom reports

Purpose:

To configure SSB to create custom reports, complete the following steps with a user that has read & write/perform access to the **use static subchapters** privilege.

Steps:

Step 1. Login to the SSB web interface, and navigate to **Reports > Configuration**.

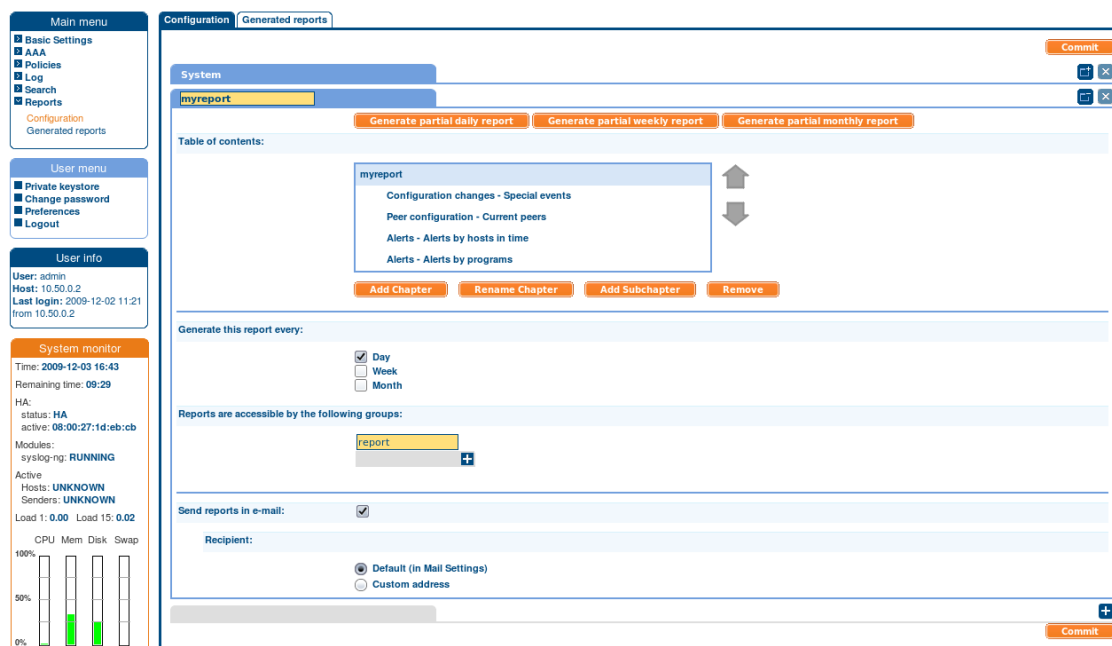


Figure 12.15. Configuring custom reports

- Step 2. Click **+** and enter a name for the custom report.
- Step 3. Reports are organized into chapters and subchapters. Select **Table of contents** > **Add Chapter**, enter a name for the chapter, then click **OK**. Repeat this step to create further chapters if needed.
- Step 4. Select **Add Subchapter** to add various reports and statistics to the chapter. The available reports will be displayed in a popup window. The reports created from custom statistics are listed at the end.
- Step 5. Use the arrows to change the order of the subchapters if needed.
- Step 6. Select how often shall SSB create the report from the **Generate this report every** field. Weekly reports are created on Mondays, while monthly reports on the first day of the month. If you want to generate the report only manually, leave this field empty.
- Step 7. By default, members of the *search* group can access the custom reports via the SSB web interface. To change this, enter the name of a different group into the **Reports are accessible by the following groups** field, or click **+** to grant access to other groups.

**Note**

Members of the listed groups will be able to access only these custom reports even if their groups does not have read access to the **Reporting** > **Reports** page. However, only those reports will be listed, to which their group has access to.

- Step 8. By default, SSB sends out the reports in e-mail to the address set in the **Basic Settings** > **Management** > **Mail settings** > **Send reports to** field.

**Note**

If this address is not set, the report is sent to the SSB administrator's e-mail address.

- To disable e-mail sending, unselect the **Send reports in e-mail** option.
- To receive e-mails only when at least one audit trail matching the search criteria was found, unselect the **Send even empty reports** option.
- To e-mail the reports to a different address, select **Recipient > Custom address**, and enter the e-mail address where the reports should be sent. Click **+** to list multiple e-mail addresses if needed.

Step 9. Click **Commit**.



Chapter 13. Classifying messages with pattern databases

Using the pattern database allows you to classify messages into various categories, receive alerts on certain messages, and to collect unknown messages using artificial ignorance.



Note

Note that the classification of messages is always performed; but its results are used only if you specifically enable the relevant options on the **Log > Options** page.

The screenshot shows the 'Log > Options' page. The left sidebar contains a 'Main menu' with items like Basic Settings, AAA, Policies, Log, Sources, Spaces, Destinations, Paths, Options (highlighted), Pattern Database, Search, and Reports. Below it is a 'User menu' with Private keystore, Change password, Preferences, and Logout. At the bottom is 'User info' showing the user is 'admin' on host '10.50.0.6' with a last login on '2010-06-17 14:17'. The main content area has tabs for Sources, Spaces, Destinations, Paths, Options (selected), and Pattern Database. The 'Options' section includes 'Options', 'TLS settings', 'SNMP source' (checked), 'Alerting' (checked), 'Rate limit' (10 alerts/minute), 'Alert email' (empty field), 'Artificial Ignorance' (checked), 'Database retention' (7 days), 'Report email' (empty field), 'Name resolving', and 'Dashboard statistics' (checked). A 'Commit' button is at the bottom right.

Figure 13.1. Enabling artificial ignorance and pattern-matching alerts

- To receive alerts on messages classified as Violation, navigate to **Log > Options** and enable the **Alerts** option.
- To receive reports on messages not included in the pattern database, navigate to **Log > Options** and enable the **Artificial ignorance** option.

13.1. The structure of the pattern database

The pattern database is organized as follows:

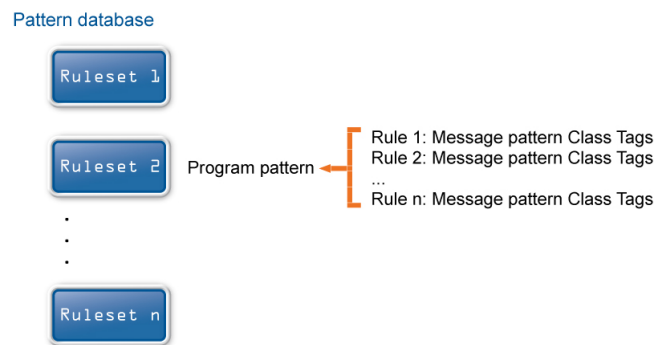


Figure 13.2. The structure of the pattern database

- The pattern database consists of rulesets. A ruleset consists of a Program Pattern and a set of rules: the rules of a ruleset are applied to log messages if the name of the application that sent the message matches the Program Pattern of the ruleset. The name of the application (the content of the \$PROGRAM macro) is compared to the Program Patterns of the available rulesets, and then the rules of the matching rulesets are applied to the message.
- The Program Pattern can be a string that specifies the name of the application or the beginning of its name (for example, to match for sendmail, the program pattern can be sendmail, or just send), and the Program Pattern can contain pattern parsers. Note that pattern parsers are completely independent from the syslog-ng parsers used to segment messages. Additionally, every rule has a unique identifier: if a message matches a rule, the identifier of the rule is stored together with the message.
- Rules consist of a message pattern and a class. The Message Pattern is similar to the Program Pattern, but is applied to the message part of the log message (the content of the \$MESSAGE macro). If a message pattern matches the message, the class of the rule is assigned to the message (for example, Security, Violation, and so on).
- Rules can also contain additional information about the matching messages, such as the description of the rule, an URL, name-value pairs, or free-form tags. This information is displayed by the syslog-ng Store Box in the e-mail alerts (if alerts are requested for the rule), and are also displayed on the search interface.
- Patterns can consist of literals (keywords, or rather, keycharacters) and pattern parsers.



Note

If the \$PROGRAM part of a message is empty, rules with an empty Program Pattern are used to classify the message.

If the same Program Pattern is used in multiple rulesets, the rules of these rulesets are merged, and every rule is used to classify the message. Note that message patterns must be unique within the merged rulesets, but the currently only one ruleset is checked for uniqueness.



13.2. How pattern matching works

A sample log message:




Figure 13.3. Applying patterns

The followings describe how patterns work. This information applies to program patterns and message patterns alike, even though message patterns are used to illustrate the procedure.

Patterns can consist of literals (keywords, or rather, keycharacters) and pattern parsers. Pattern parsers attempt to parse a sequence of characters according to certain rules.



Note

Wildcards and regular expressions cannot be used in patterns. The @ character must be escaped, that is, to match for this character, you have to write @@ in your pattern. This is required because pattern parsers of syslog-ng are enclosed between @ characters.

When a new message arrives, syslog-ng attempts to classify it using the pattern database. The available patterns are organized alphabetically into a tree, and syslog-ng inspects the message character-by-character, starting from the beginning. This approach ensures that only a small subset of the rules must be evaluated at any given step, resulting in high processing speed. Note that the speed of classifying messages is practically independent from the total number of rules.

For example, if the message begins with the *Apple* string, only patterns beginning with the character *A* are considered. In the next step, syslog-ng selects the patterns that start with *Ap*, and so on, until there is no more specific pattern left.

Note that literal matches take precedence over pattern parser matches: if at a step there is a pattern that matches the next character with a literal, and another pattern that would match it with a parser, the pattern with the literal match is selected. Using the previous example, if at the third step there is the literal pattern *Apport* and a pattern parser *Ap@STRING@*, the *Apport* pattern is matched. If the literal does not match the incoming string (for example, *Apple*), syslog-ng attempts to match the pattern with the parser. However, if there are two or more parsers on the same level, only the first one will be applied, even if it does not perfectly match the message.

If there are two parsers at the same level (for example, *Ap@STRING@* and *Ap@QSTRING@*), it is random which pattern is applied (technically, the one that is loaded first). However, if the selected parser cannot parse at least one character of the message, the other parser is used. But having two different parsers at the same level is extremely rare, so the impact of this limitation is much less than it appears.

13.3. Searching for rulesets

To display the rules of a ruleset, enter the name of the ruleset into the **Search > Ruleset name** field, and click **Show**. If you do not know the name of the ruleset, type the beginning letter(s) of the name, and the names of the matching rulesets will be displayed. If you are looking for a specific rule, enter a search term into the **Program** or **Message** field and select **Search**. The rulesets that contain matching rules will be displayed.

**Note**

Rulesets containing large number of rules may not display correctly.

Database

Publication date: 2008-10-29 (version 2)
Rulesets/rules: 280/6007

Upload database: Choose... Export Upload

Search

Ruleset name: Show

Match for log message: Search

Create new ruleset

hddtemp

Name:

URL:

Program pattern:

Description:

Rules:

Pattern	Class
/dev@QSTRING:::@@QSTRING:::@ @NUMBER:::@	System
/dev/sg@NUMBER:::@@QSTRING:::@ drive is sleeping	System
/dev@QSTRING:::@@QSTRING:::@ drive is sleeping	System
/dev/sg@NUMBER:::@@QSTRING:::@ @NUMBER:::@	System

Export ruleset Delete Commit

System monitor

Time: 2009-12-03 16:43
Remaining time: 09:29

HA:
status: HA
active: 08:00:27:1d:eb:cb

Modules:
syslog-ng: RUNNING

Active
Hosts: UNKNOWN
Senders: UNKNOWN

Load 1: 0.00 Load 15: 0.02

CPU Mem Disk Swap

100%
50%
0%

1% 34% 25% 0%

Figure 13.4. Searching rules

13.4. Procedure – Creating new rulesets and rules

Purpose:

To create a new ruleset and new rules, complete the following steps:

Steps:

Step 1. Select **Log > Pattern Database > Create new ruleset**.

**Tip**

If you search for a ruleset that does not exist, SSB offers you to create a new ruleset with the name you were searching for.

Step 2. Enter a name for the ruleset into the **Name** field.



The screenshot displays the 'New Ruleset' configuration page in the Syslog-ng Premium Edition web interface. On the left, a sidebar contains navigation links: Search, Reports, User menu (Private key store, Change password, Preferences, Logout), User info (User: admin, Host: 10.50.0.2), and System monitor (Time: 2010-06-17 13:46, Remaining time: 09:35, Locked: admin@10.60.0.2, Modules: syslog-ng: RUNNING, Active Hosts: 0, Senders: 0, Load 1: 0.01, Load 15: 0.07, CPU: 1%, Mem: 31%, Disk: 7%, Swap: 0%).

The main form is titled 'New Ruleset' and contains the following sections:

- Search:** Fields for 'Ruleset name' and 'Match for log message' (Program and Message), with 'Show' and 'Search' buttons.
- Create new ruleset:** A button to initiate the process.
- Form Fields:**
 - Name:** 'New Ruleset'
 - URL:** (empty)
 - Program pattern:** 'myapplication'
 - Description:** (empty text area)
- Rules:**
 - Pattern:** 'User @stringusername@ successfully logged in'
 - Class:** 'System' (selected from a dropdown)
 - Description:** (empty text area)
 - URL:** (empty)
 - Tags:** 'login' (added from a list)
 - Rule ID:** (empty)
 - Values:**

Name	Value
username	\$username
eventtype	login

A 'Commit' button is located at the bottom right of the form.

Figure 13.5. Creating pattern database rulesets

- Step 3. Enter the name of the application or a pattern that matches the applications into the **Program pattern** field. For details, see *Section 13.7, Using pattern parsers (p. 188)*.
- Step 4. Optionally, add a description to the ruleset.
- Step 5. Add rules to the class.
- Step a. Click **+** in the **Rules** section.
- Step b. Enter the beginning of the log message or a pattern that matches the log message into the **Pattern** field. For details, see *Section 13.7, Using pattern parsers (p. 188)*. Note that only messages sent by applications matching the **Program pattern** will be affected by this pattern.
- Step c. Select the type of the message from the **Class** field. This class will be assigned to messages matching the pattern of this rule. The following classes are available: *Violation*, *Security*, and *System*.
If alerting is enabled at **Log > Options > Alerting**, SSB automatically sends an alert if a message is classified as Violation.
- Step d. Optionally, you can add a description, custom tags, and name-value pairs to the rule. Note that the values of name-value pairs can contain macros in the `${macroname}` format. For details on pattern databases and macros, see *The syslog-ng Premium Edition Administrator Guide*, available at the [BalaBit Documentation Page](#).



Step 6. Repeat the previous step to add more rules.

Step 7. Click **Commit**.

13.5. Exporting databases and rulesets

To export the entire pattern database, navigate to **Log > Pattern Database** and select **Export**.

To export a ruleset, enter the name of the ruleset into the **Search > Ruleset name** field, click **Show**, and select **Export ruleset**. If you do not know the name of the ruleset, enter a search term into the **Program** or **Message** field and select **Search**. The rulesets that contain matching rules will be displayed.

13.6. Importing pattern databases

You can upload official databases distributed by BalaBit or pattern databases that you have exported from SSB. Official databases are available at the BalaBit website. The official database currently includes patterns for various Linux/Unix applications, and is essentially a conversion of the database used by the [logcheck application](#). Other databases for Cisco devices and Microsoft Windows applications are under development.

To import a ruleset, navigate to **Log > Pattern Database** and select **Browse**. Then locate the database file to upload, and click **Upload**.



Note

Imported rules are effective immediately after the upload is finished.

If you have modified a rule that was originally part of an official database, then the update will not modify this rule.

13.7. Using pattern parsers

Pattern parsers attempt to parse a part of the message using rules specific to the type of the parser. Parsers are enclosed between @ characters. The syntax of parsers is the following:

- a beginning @ character;
- the type of the parser written in capitals;
- optionally a name;
- parameters of the parser, if any;
- a closing @ character.



Example 13.1. Pattern parser syntax

A simple parser:

```
@STRING@
```

A named parser:

```
@STRING:myparser_name@
```

A named parser with a parameter:

```
@STRING:myparser_name: *@
```

A parser with a parameter, but without a name:



```
@STRING::*@
```

The following parsers are available:

- **@ANYSTRING@**: Parses everything to the end of the message; you can use it to collect everything that is not parsed specifically to a single macro. In that sense its behavior is similar to the *greedy()* option of the CSV parser.
- **@DOUBLE@**: An obsolete alias of the **@FLOAT@** parser.
- **@ESTRING@**: This parser has a required parameter that acts as the stopcharacter: the parser parses everything until it find the stopcharacter. For example to stop by the next " (double quote) character, use **@ESTRING::"@**. As of syslog-ng 3.1, it is possible to specify a stopstring instead of a single character, for example **@ESTRING::stop_here.@"**.
- **@FLOAT@**: A floating-point number that may contain a dot (.) character. (Up to syslog-ng 3.1, the name of this parser was **@DOUBLE@**.)
- **@IPv4@**: Parses an IPv4 IP address (numbers separated with a maximum of 3 dots).
- **@IPv6@**: Parses any valid IPv6 IP address.
- **@IPVANY@**: Parses any IP address.
- **@NUMBER@**: A sequence of decimal (0-9) numbers (for example 1, 0687, and so on). Note that if the number starts with the 0x characters, it is parsed as a hexadecimal number, but only if at least one valid character follows 0x.
- **@QSTRING@**: Parse a string between the quote characters specified as parameter. Note that the quote character can be different at the beginning and the end of the quote, for example: **@QSTRING::"@** parses everything between two quotation marks ("), while **@QSTRING::<>@"** parses from an opening bracket to the closing bracket.
- **@STRING@**: A sequence of alphanumeric characters (0-9, A-z), not including any whitespace. Optionally, other accepted characters can be listed as parameters (for example to parse a complete sentence, add the whitespace as parameter, like: **@STRING:: @**). Note that the @ character cannot be a parameter, nor can line-breaks or tabs.

Patterns and literals can be mixed together. For example, to parse a message that begins with the *Host:* string followed by an IP address (for example *Host: 192.168.1.1*), the following pattern can be used: **Host:@IPv4@**.



Note

Note that using parsers is a CPU-intensive operation. Use the **ESTRING** and **QSTRING** parsers whenever possible, as these can be processed much faster than the other parsers.



Example 13.2. Using the **STRING** and **ESTRING** parsers

For example, if the message is *user=joe96 group=somegroup, @STRING:mytext:@* parses only to the first non-alphanumeric character (=), parsing only *user=joe96*. **@STRING:mytext:=@** parses the equation mark as well, and proceeds to the next non-alphanumeric character (the whitespace), resulting in *user=joe96* being parsed. **@STRING:mytext:=@** will parse the whitespace as well, and proceed to the next non-alphanumeric non-equation mark non-whitespace character, resulting in *user=joe96 group=somegroup*.

Of course, usually it is better to parse the different values separately, like this: *"user=@STRING:user@ group=@STRING:group@"*.



If the username or the group may contain non-alphanumeric characters, you can either include these in the second parameter of the parser (as shown at the beginning of this example), or use an ESTRING parser to parse the message till the next whitespace: `"user=@ESTRING:user: @group=@ESTRING:group: @"`.

13.8. Procedure – Using parser results in filters and templates

Purpose:

The results of message classification and parsing can be used in custom filters and file and database templates as well. There are two built-in macros in SSB that allow you to use the results of the classification: the `.classifier.class` macro contains the class assigned to the message (for example violation, security, or unknown), while the `.classifier.rule_id` macro contains the identifier of the message pattern that matched the message.



Note

ID of the message pattern is automatically inserted into the template if the messages are forwarded to an SQL database.

To use these macros as filters in a log path, complete the following procedure:

Steps:

Step 1. Navigate to **Log > Paths** and select the log path to use.

Step 2.

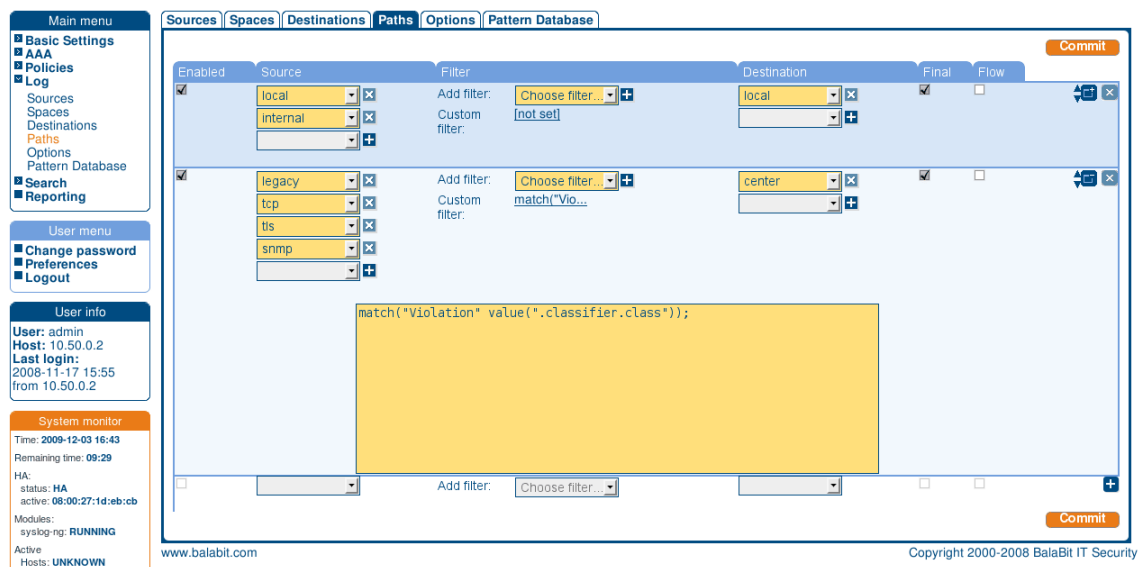


Figure 13.6. Filtering messages based on the classification

To filter on a specific message class, select **Add filter > classifier_class**, select **+**, then select the class to match (for example *Violation*) from the **classifier_class** field.



Step 3. To filter on messages matching a specific classification rule, **Add filter** > **classifier_rule_id**, select **+**, then enter the unique identifier of the rule (for example `e1e9c0d8-13bb-11de-8293-000c2922ed0a`) into the **classifier_rule_id** field.



Note

To filter messages based on other classification data like tags, you have to use Custom filters. For details, see *Section 10.3, Filtering messages (p. 146)*.

Step 4. Click **Commit**.

13.9. Using the values of pattern parsers in filters and templates

Similarly, to *Procedure 13.8, Using parser results in filters and templates (p. 190)*, the results of pattern parsers can be used as well. To accomplish this, you have to add a name to the parser, and then you can use this name as a macro that refers to the parsed value of the message.

For example, you want to parse messages of an application that look like *"Transaction: <type>."*, where *<type>* is a string that has different values (for example refused, accepted, incomplete, and so on). To parse these messages, you can use the following pattern:

```
'Transaction: @ESTRING:..@'
```

Here the `@ESTRING@` parser parses the message until the next full stop character. To use the results in a filter or a filename template, include a name in the parser of the pattern, for example:

```
'Transaction:
  @ESTRING:TRANSACTIONTYPE:..@'
```

After that, add a custom template to the logpath that uses this template. For example, to select every *accepted* transaction, use the following custom filter in the log path:

```
match("accepted" value("TRANSACTIONTYPE"));
```



Note

The above macros can be used in database columns and filename templates as well, if you create custom templates for the destination or logspace.



Chapter 14. Troubleshooting SSB

This section describes the tools to detect networking problems, and also how to collect core files and view the system logs of SSB.

14.1. Procedure – Network troubleshooting

Purpose:

The **Troubleshooting** menu provides a number of diagnostic commands to resolve networking issues. Logfiles of SSB can also be displayed here — for details, see *Procedure 14.3, Viewing logs on SSB (p. 193)*.

Figure 14.1. Network troubleshooting with SSB

The following commands are available:

- ping: Sends a simple message to the specified host to test network connectivity.
- traceroute: Sends a simple message from SSB to the specified host and displays all hosts on the path of the message. It is used to trace the path the message travels between the hosts.



- **connect**: Attempts to connect the specified host using the specified port. It is used to test the availability or status of an application on the target host.

To execute one of the above commands, complete the following steps:

Steps:

- Step 1. Navigate to **Basic Settings > Troubleshooting**.
- Step 2. Enter the IP address or the hostname of the target host into the **Hostname** field of the respective command. For the Connect command, enter the target port into the **Port** field.
- Step 3. Click the respective action button to execute the command.
- Step 4. Check the results in the popup window. Log files are displayed in a separate browser window.

14.2. Gathering data about system problems

SSB automatically generates core files if an important software component (for example syslog-ng, or the indexer) of the system crashes for some reason. These core files can be of great help to the BalaBit Support Team to identify problems. When a core file is generated, the SSB administrator receives an alerting e-mail, and an SNMP trap is generated if alerting is properly configured (for details, see *Section 4.6, Configuring system monitoring on SSB (p. 47)* and *Section 4.5, SNMP and e-mail alerts (p. 42)*). To display a list of alerts if monitoring is not configured, navigate to **Search > Log Alerts**.

To list and download the generated core files, navigate to **Basic Settings > Troubleshooting > Core files**.

By default, core files are deleted after 14 days. To change the deletion timeframe, navigate to **Basic Settings > Management > Core files**.

14.3. Procedure – Viewing logs on SSB

Purpose:

The **Troubleshooting** menu provides an interface to view the logs generated by the various components of SSB. For details on how to browse the log messages received by SSB from its peers, see *Chapter 12, Browsing log messages and SSB reports (p. 156)*.



Note

Because of performance reasons, log files larger than 2 Megabytes are not displayed in the web interface. To access these logs, download the file instead.

Steps:

- Step 1. Navigate to **Basic Settings > Troubleshooting > View log files**.
- Step 2. Use the **Logtype** roll-down menu to select the message type.
 - **SSB**: Logs of the SSB web interface.
 - **syslog**: All system logs of the SSB host.



- *syslog-ng*: Internal log messages of the built-in syslog-ng server. These logs do not contain messages received from the peers.

Step 3.

- To download the log file, click **Download**.
- To follow the current log messages real-time, click **Tail**.
- To display the log messages, click **View**.

Step 4. To display log messages of the last seven days, select the desired day from the **Day:** field and click **View**.



Tip

To display only the messages of a selected host or process, enter the name of the host or process into the **Message:** field.

The **Message:** field acts as a generic filter: enter a keyword or a POSIX (basic) regular expression to display only messages that contain the keyword or match the expression.

14.4. Procedure – Collecting logs and system information for error reporting

Purpose:

To track down support requests, the BalaBit Support Team might request you to collect system-state and debugging information. This information is collected automatically, and contains log files, the configuration file of SSB, and various system-statistics.



Note

Sensitive data like key files and passwords are automatically removed from the files.

The **Basic Settings > Management > Debug logging > Enable debug logs** option is not related to the verbosity of log messages: it adds the commands executed by the SSB web interface to the log.

To collect system-state information, navigate to **Basic Settings > Troubleshooting > System debug** and click **Collect and save current system state info**, then save the created zip file. The name of the file uses the *debug_info-<hostname>YYYYMMDDHHMM* format.

To collect information for a specific error, complete the following steps:

Steps:

Step 1. Navigate to **Basic Settings > Troubleshooting > System debug**.



The screenshot shows the 'Troubleshooting' tab of a network management interface. On the left, there are three sidebars: 'Main menu' with links like Basic Settings, AAA, Policies, Log, Search, and Reports; 'User menu' with links like Private keystore, Change password, Preferences, and Logout; and 'User info' showing the user is 'admin' with host '10.50.0.2'. Below these is a 'System monitor' section displaying system time, remaining time, HA status, active IP, modules, active hosts, senders, and load. It also includes a bar chart for CPU (1%), Mem (34%), Disk (25%), and Swap (0%) usage.

The main content area contains several diagnostic tools:

- Ping:** Hostname: 127.0.0.1, Ping host button.
- Traceroute:** Hostname: 127.0.0.1, Traceroute host button.
- Connect to TCP port:** Hostname: 127.0.0.1, TCP port: 22, Connect to host button.
- View log files:** Logtype: syslog, Day: Friday, Search patterns and Message input fields, Download, View, and Tail buttons.
- System debug:** A button to 'Collect and save current system state in', a status indicator 'Debug mode is OFF', Start and Stop buttons, and a 'Save collected debug info' button.
- Core files:** A table with headers: File, Size, Creation date, and Encrypt for download.

Figure 14.2. Collecting debug information

Step 2. Click **Start**.



Note

Starting debug mode increases the log level of SSB, and might cause performance problems if the system is under a high load.

Step 3. Reproduce the event that causes the error, for example send a log message from a client.

Step 4. Click **Stop**.

Step 5. Click **Save the collected debug info** and save the created zip file. The name of the file uses the *debug_info-<hostname>YYYYMMDDHHMM* format.

Step 6. Attach the file to your support ticket.

14.5. Status history and statistics

SSB displays various statistics and status history of system data and performance on the dashboard at **Basic Settings > Dashboard**. The dashboard is essentially an extension of the system monitor: the system monitor displays only the current values, while the dashboard creates graphs and statistics of the system parameters.

The dashboard consists of different modules. Every module displays the history of a system parameter for the current day. To display the graph for a longer period (last week, last month, or last year), select the **Week**, **Month**, or **Year** options, respectively. Hovering the mouse over a module enlarges the graph and displays the color code used on the graph.

To display statistics of a module as a table for the selected period, click on the graph.

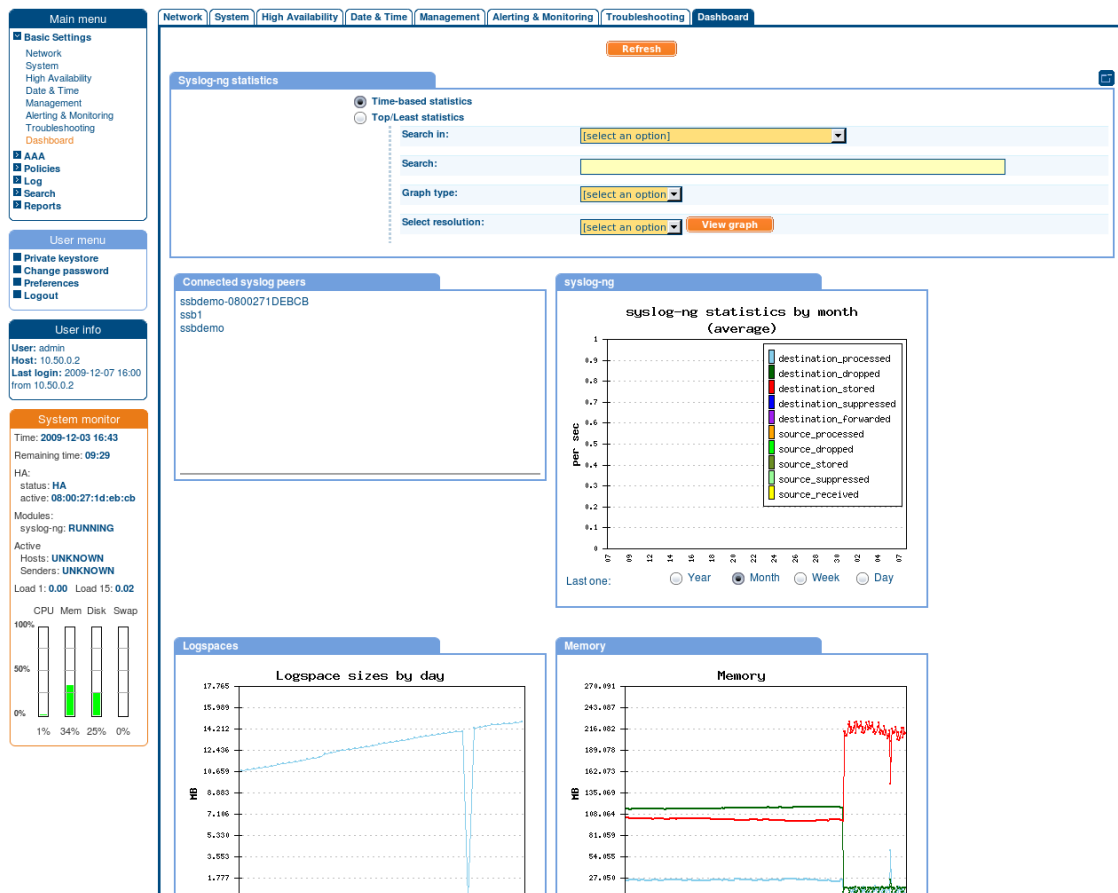


Figure 14.3. The dashboard

The following modules are displayed on the dashboard of SSB:

- **syslog-ng:** syslog-ng statistics about the received, processed, and dropped messages. See also *Procedure 14.5.1, Displaying custom syslog-ng statistics (p. 197)*.
- **Connected syslog peers:** A list of hosts that actively send messages to SSB. Note that these values are updated periodically based on the **Sampling interval** set on page **Log > Options > Dashboard Statistics**. For details, see *Procedure 14.5.1, Displaying custom syslog-ng statistics (p. 197)*.



- **syslog-ng statistics:** The rate of incoming messages in messages/second. Note that the values displayed are an average values calculated for the last fifteen minutes.
- **Logspaces:** The size of the logspaces. Note that these values are updated only in every ten minutes.
- **Memory:** The memory used by the system.
- **Disk:** Filesystem usage for the different partitions.
- **CPU:** CPU usage.
- **Network connections:** Number of network connections.
- **External interface:** Traffic on the external interface.
- **Management interface:** Traffic on the management interface.
- **Load average:** Average load of the system.
- **Processes:** The number of running processes.

For details about setting the statistics collection options, see *Section 12.7, Statistics collection options (p. 177)*

14.5.1. Procedure – Displaying custom syslog-ng statistics

Purpose:

To display statistics of a specific source, destination, or host, complete the following procedure:

Steps:

- Step 1. Navigate to **Basic Settings > Dashboard > syslog-ng statistics**.
- Step 2.
 - To display the statistics of a destination file, select *destination* from the **Search in** field, and enter the name of the destination into the **Search** field. Destinations name all start with the *ds* characters.
 - To display the statistics of a particular host, select *source* from the **Search in** field, and enter the hostname or IP address of the host into the **Search** field.
- Step 3. Select the time period to display from the **Select resolution** field.
- Step 4. Click **View graph**.

14.6. Troubleshooting an SSB cluster

The following sections help you to solve problems related to high availability clusters.

- For details on how to recover a cluster that has broken down, see *Procedure 14.6.1, Recovering SSB if both nodes broke down (p. 197)*.
- For details on how to resolve a split-bran situation when the nodes of the cluster were simultaneously active for a time, see *Procedure 14.6.2, Recovering from a split brain situation (p. 198)*.
- For details on replacing a broken node with a new appliance, see *Procedure 14.6.3, Replacing a node in an SSB HA cluster (p. 200)*.

14.6.1. Procedure – Recovering SSB if both nodes broke down

Purpose:



It can happen that both nodes break down simultaneously (for example because of a power failure), or the slave node breaks down before the original master node recovers. To properly recover SSB, complete the following steps:



Note

As of SSB version 1.1.1, when both nodes of a cluster boot up in parallel, the node with the 1.2.4.1 HA IP address will become the master node.

Steps:

Step 1. Power off both nodes by pressing and releasing the power button.



Warning

Hazard of data loss! If SSB does not shut off, press and hold the power button for approximately 4 seconds. This method terminates connections passing SSB and might result in data loss.

Step 2. Power on the node that was the master before SSB broke down. Consult the system logs to find out which node was the master before the incident: when a node boots as master, or when a takeover occurs, SSB sends a log message identifying the master node.



Tip

Configure remote logging to send the log messages of SSB to a remote server where the messages are available even if the logs stored on SSB become inaccessible. For details on configuring remote logging, see *Section 4.5, SNMP and e-mail alerts (p. 42)*.

Step 3. Wait until this node finishes the boot process.

Step 4. Power on the other node.

14.6.2. Procedure – Recovering from a split brain situation

Purpose:

A split brain situation is caused by a temporary failure of the network link between the cluster nodes, resulting in both nodes switching to the active (master) role while disconnected. This might cause that new data (for example log messages) is created on both nodes without being replicated to the other node. Thus, it is likely in this situation that two diverging sets of data have been created, which cannot be trivially merged.



Warning

Hazard of data loss! In a split brain situation, valuable log messages might be available on both SSB nodes, so special care must be taken to avoid data loss.



The nodes of the SSB cluster automatically recognize the split brain situation once the connection between the nodes is reestablished, and do not perform any data synchronization to prevent data loss. When a split brain situation is detected, it is visible on the SSB system monitor, in the system logs (*Split-Brain detected, dropping connection!*), and SSB sends an alert as well.

To recover an SSB cluster from a split brain situation, complete the following steps.



Warning
Do NOT shut down the nodes.

Steps:

Step 1. Temporarily disable all incoming traffic. Navigate to **Basic Settings > System > Traffic control** and click **Disable**.

If the web interface is not accessible or unstable, complete the following steps:

- Step a. Login to SSB as *root* locally (or remotely using SSH) to access the Console menu.
- Step b. Select **Shells > Core Shell**, and issue the `syslog-ng stop` command.
- Step c. Issue the `date` and check the system date and time. If it is incorrect (for example it displays 2000 January), replace the system battery. For details, see the hardware manual of the appliance.
- Step d. Repeat the above steps on the other SSB node.

Step 2. *Optional step for data recovery:* Check the log spaces saved on the SSB nodes.

- Step a. Login to the node from a local console.
- Step b. Select **Shells > Core Shell** and enter `cd /opt/ssb/var/logspace/`. The log spaces are located under this directory.
- Step c. Find which files were modified since the split brain situation occurred. Use the `find . -mtime -n` to find the files modified during the last $n \times 24$ hours, or the `find . -mmin -n` to find the files modified during the last n minutes.

Step 3. Decide which node should be the master node from now on, then perform the following steps on the to-be-slave node:

- Step a. Login to the node from a local console.
- Step b. *Optional step for data recovery:* Backup the log messages that were modified since the split brain situation occurred.



Warning
This data will be deleted from the SSB node when the split-brain situation is resolved. There is no way to import this data back into the database of SSB; it will be available only for offline use.



Step c. *Optional step for data recovery:* Type `exit` to return to the console menu.

Step d. Select **Shells** > **Boot shell**. If the to-be-slave node is not already the slave node, fail over the cluster to the other node manually by issuing the `/usr/share/heartbeat/hb_standby` command.

Step e. Stop the core firmware. Issue the `/etc/init.d/boot-xcb stop` command.

Step f. Invalidate the DRBD. Issue the following commands:

```
/sbin/drbdsetup /dev/drbd0 disconnect
```

```
/sbin/drbdsetup /dev/drbd0 invalidate.
```

Step 4. Reboot the to-be-slave node.

Step 5. Reboot the to-be-master node. The SSB cluster will be now functional, accepting traffic as before.

Step 6. After both nodes reboot, the cluster should be in **Degraded Sync** state, the master being **SyncSource** and the slave being **SyncTarget**. The master node should start synchronizing its data to the slave node. Depending on the amount of data, this can take a long time. To adjust the speed of the synchronization, see *Section 6.2.2, Adjusting the synchronization speed of DRBD (p. 88)*.

14.6.3. Procedure – Replacing a node in an SSB HA cluster

Purpose:

To replace a unit in an SSB cluster with a new appliance, complete the following steps.

Steps:

Step 1. Verify the HA status on the working node. Select **Basic Settings** > **High Availability**. If one of the nodes has broken down or is missing, the **Status** field displays *DEGRADED*.

Step 2. Note down the IP addresses of the **Heartbeat** and the **Next hop monitoring** interfaces.

Step 3. Perform a full system backup. Before replacing the node, create a complete system backup of the working node. For details, see *Section 4.7, Data and configuration archiving and backups (p. 54)*.

Step 4. Check which firmware version is running on the working node. Select **Basic Settings** > **System** > **Version details** and write down the exact version numbers.


Step 5. Login to your MyBalaBit account at <https://www.balabit.com/mybalabit/> and download the CD ISO for the same SSB version that is running on your working node.

Step 6. Without connecting the replacement unit to the network, install the replacement unit from the ISO file. Use the IPMI interface if needed.

Step 7. When the installation is finished, connect the two SSB units with an Ethernet cable via the Ethernet connectors labeled as 4 (or *HA*).

Step 8. Reboot the replacement unit and wait until it finishes booting.

Step 9. Login to the working node and verify the HA state. Select **Basic Settings** > **High Availability**. The **Status** field should display *HALF*.

Step 10. Reconfigure the IP addresses of the **Heartbeat** and the **Next hop monitoring** interfaces. Click .

Step 11. Click **Other node** > **Join HA**.



Step 12. Click **Other node > Reboot**.

Step 13. The replacement unit will reboot and start synchronizing data from the working node. The **Basic Settings > High Availability > Status** field will display *DEGRADED SYNC* until the synchronization finishes. Depending on the size of the hard disks and the amount of data stored, this can take several hours.

Step 14. After the synchronization is finished, connect the other Ethernet cables to their respective interfaces (external to 1 or *EXT*, management to 2 or *MGMT*) as needed for your environment.

Expected result:

A node of the SSB cluster is replaced with a new appliance.

14.7. Procedure – Restoring SSB configuration and data

Purpose:

The following procedure describes how to restore the configuration and data of SSB from a complete backup, for example, after a hardware replacement.

Steps:

Step 1. Connect to your backup server and locate the directory where SSB saves the backups. The configuration backups are stored in the *config* subdirectory in timestamped files. Find the latest configuration file (the configuration files are called *SSB-timestamp.config*).

Step 2. Connect to SSB.

If you have not yet completed the Welcome Wizard, click **Browse**, select the configuration file, and click **Import**.

If you have already completed the Welcome Wizard, navigate to **Basic Settings > System > Import configuration > Browse**, select the configuration file, and click **Import**.

Step 3. Navigate to **Policies > Backup & Archive/Cleanup**. Verify that the settings of the target servers and the backup protocols are correct.

Step 4. Navigate to **Basic Settings > Management > System backup**, click **Restore now** and wait for the process to finish. Depending on the amount of data stored in the backup, and the speed of the connection to the backup server, this may take a long time.

Step 5. Navigate to **Log > Spaces**, and click **Restore ALL**. Depending on the amount of data stored in the backup, and the speed of the connection to the backup server, this may take a long time.

14.8. Procedure – SAN troubleshooting

Purpose:

When experiencing unexpected behavior regarding iSCSI storage, or when one of the following issues arise, perform the following steps.

- Volumes are not visible on the **Basic Settings > Storage** page, and they do not reappear when clicking **Rescan**.
- "Missing volumes" messages appear on the **Log > Spaces** page.



- When adding volumes to logspaces, error messages appear during commit.
- During the boot process, the console prints the following messages (visible in dmesg):
device-mapper: table: 253:0: multipath: error getting device or device-mapper:
ioctl: error adding target to table.

Prerequisites:

Console or ILOM access to the machine. The procedure itself requires approximately 5 minutes to complete.

Steps:

- Step 1. Reboot the system through console or ILOM. When the **QLogic Fast!UTIL** boot screen appears, press **CTRL+Q** to display the iSCSI adapter configuration tool.
- Step 2. Select **Configuration Settings** and press **Enter**.
- Step 3. Select **Restore Adapter Defaults** and press **Enter**.
- Step 4. When the operation has finished, press any key to return to the **Configuration Settings** menu.
- Step 5. Select **Clear Persistent Targets** and press **Enter**.
- Step 6. When the operation has finished, press any key to return to the **Configuration Settings** menu.
- Step 7. Press **ESC** to return to the main menu.
- Step 8. Select **Exit Fast!UTIL** and press **Enter**.
- Step 9. Select **Reboot System** and press **Enter** to reboot the machine.



Appendix A. Package contents inventory

Carefully unpack all server components from the packing cartons. The following items should be packaged with the syslog-ng Store Box:

- A syslog-ng Store Box appliance, preinstalled with the latest syslog-ng Store Box firmware.
- syslog-ng Store Box accessory kit, including the following:
 - Delivery note
 - syslog-ng Store Box License Agreement
 - syslog-ng Store Box Certificate (includes the purchased license and support options, and support contact details)



Note

The default BIOS and IPMI passwords are in the documentation.

- syslog-ng Store Box Hardware Installation Guide
- Rackmount
- Power cable



Appendix B. syslog-ng Store Box Hardware Installation Guide

This leaflet describes how to set up the syslog-ng Store Box (SSB) hardware. Refer to the following documents for step-by-step instructions:

- *syslog-ng Store Box SSB1000*: For details on installing SSB into a rack, see the *SC811 CHASSIS Series User's Manual, Chapter 6: Rack Installation*. For details on connecting the cables to SSB see the *X8SIE/X8SIE-F/X8SIE-LN4/X8SI6-F User's Manual, Section 2-5 Connectors/IO Ports*.
- *syslog-ng Store Box SSB1000d and SAN Connect*: For details on installing SSB into a rack, see the *SC825M CHASSIS Series User's Manual, Chapter 5: Rack Installation*. For details on connecting the cables to SSB see the *X8DT3/X8DTi/X8DT3-F/X8DTi-F/X8DT3-LN4F/X8DTi-LN4F User's Manual, Section 2-5 Control Panel Connectors/IO Ports*.
- *syslog-ng Store Box SSB5000, and SSB10000* : For details on installing SSB into a rack, see the *SC826 CHASSIS Series User's Manual, Chapter 6: Rack Installation*. For details on connecting the cables to SSB see the *X8DT3/X8DTi/X8DT3-F/X8DTi-F/X8DT3-LN4F/X8DTi-LN4F User's Manual, Section 2-5 Control Panel Connectors/IO Ports*.

The manuals are available online at [the BalaBit Documentation page](#).

- For details on how to install a single SSB unit, see *Procedure B.1, Installing the SSB hardware (p. 204)*.
- For details on how to install a two SSB units in high availability mode, see *Procedure B.2, Installing two SSB units in HA mode (p. 205)*.
- If you have an external storage module, and you have purchased the SSB SAN Connect edition, see *Procedure B.3, Installing a SAN storage module to SSB (p. 206)*.

B.1. Procedure – Installing the SSB hardware

Purpose:

To install a single SSB unit, complete the following steps.

Steps:

Step 1. Unpack SSB.

Step 2. *Optional step*: Install SSB into a rack with the slide rails. Slide rails are available for all SSB appliances.

Step 3. Connect the cables.

Step a. Connect the Ethernet cable facing your LAN to the Ethernet connector labeled as *EXT*. This is the external interface of SSB and is used to configure SSB. (For details on the roles of the different interfaces, see *Section 2.5, Network interfaces (p. 8)*.)

Step b. Connect an Ethernet cable that you can use to remotely support the SSB hardware to the *IPMI* interface of SSB. For details, see the following documents:

- *The Onboard BMC/IPMI User's Guide*, available at [the BalaBit Hardware Documentation page](#).

**Warning**

Connect the IPMI before plugging in the power chord. Failing to do so will result in IPMI failure.

It is not necessary for the IPMI interface to be accessible from the Internet, but the administrator of SSB must be able to access it for support and troubleshooting purposes in case vendor support is needed.

**Warning**

Access to information available only via the IPMI interface is a not mandatory, but highly recommended to speed up the support and troubleshooting processes.

Step c. *Optional step:* Connect the Ethernet cable to be used for managing SSB after its initial configuration to the Ethernet connector labeled as *MGMT*. This is the management interface of SSB. (For details on the roles of the different interfaces, see *Section 2.5, Network interfaces (p. 8).*)

Step d. *Optional step:* Connect the Ethernet cable connecting SSB to another SSB node to the Ethernet connector labeled as *HA*. This is the high availability (HA) interface of SSB. (For details on the roles of the different interfaces, see *Section 2.5, Network interfaces (p. 8).*)

Step 4. Power on the hardware.

Step 5. Change the BIOS and IPMI passwords on the syslog-ng Store Box. The default password is *changeme*.

Step 6. Connect to the SSB web interface from a client machine and complete the Welcome Wizard. This might require you to configure an alias interface on the client machine. Step 5 is described in detail in *Chapter 3, The Welcome Wizard and the first login (p. 17).*

**Note**

The syslog-ng Store Box Administrator Guide is available on the SSB on the [BalaBit Documentation page](#).

B.2. Procedure – Installing two SSB units in HA mode

Purpose:

To install SSB with high availability support, complete the following steps.

Steps:

Step 1. For the first SSB unit, complete *Procedure B.1, Installing the SSB hardware (p. 204)*.

Step 2. For the second SSB unit, complete Steps 1-2 of *Procedure B.1, Installing the SSB hardware (p. 204)*.

Step 3. Connect the two units with an Ethernet cable via the Ethernet connectors labeled as *HA*.



- Step 4. Power on the second unit.
- Step 5. Connect to the SSB web interface of the first unit from a client machine and enable the high availability mode. Navigate to **Basic Settings > High Availability** . Click **Convert to Cluster**, then reload the page in your browser.
- Step 6. Reboot the slave unit, then reboot the master unit. Wait a few minutes until the master unit boots, then reload the page in your browser.
- Step 7. Wait until the slave unit synchronizes its disk to the master unit. Depending on the size of the hard disks, this may take several hours. You can increase the speed of the synchronization via the SSB web interface at **Basic Settings > High Availability > DRBD sync rate limit**.

B.3. Procedure – Installing a SAN storage module to SSB

Purpose:

If you have an external storage module, and you have purchased the SSB SAN Connect edition, complete the following steps.

Steps:

- Step 1. Install the SSB unit(s) as described in *Procedure B.1, Installing the SSB hardware (p. 204)*. If you have purchased two units for high availability, complete *Procedure B.2, Installing two SSB units in HA mode (p. 205)*.
- Step 2. Connect iSCSI port of SSB to the storage module with a crosslink cable.



Warning

If you have two SSB SAN Connect units in high availability mode, and your storage module has multiple iSCSI controllers, connect both units to the same controller, otherwise the second SSB unit might not be able to access the storage module after a failover.

- Step 3. Power on the storage module hardware.
- Step 4. Configure the storage module to accept connections from SSB. Refer the manual of your storage module for details.



Warning

Ensure that the LUN numbers assigned to the volume mappings of the storage are between 1 and 30 (including 1 and 30). Do NOT use the 0 or the 31 LUN, because SSB will not be able to access the volumes.

- Step 5. Power on the SSB hardware.
- Step 6. Configure SSB to access the storage module. For details, see *Procedure 3.3, Configuring storage access in SSB (p. 31)* and *Section 6.11, Managing SAN access in SSB (p. 112)*. For details on SAN troubleshooting, see *Procedure 14.8, SAN troubleshooting (p. 201)*.



B.4. Hardware Troubleshooting

Refer to the *Setup Troubleshooting* chapter of the respective guide if you encounter any problems. If you still experience problems, contact the BalaBit Support Team via phone or e-mail:

To access the BalaBit Online Support System (BOSS), sign up for an account at [the MyBalaBit page](#) and [request access to the BalaBit Online Support System \(BOSS\)](#). Online support is available 24 hours a day.

BOSS is available only for registered users with a valid support package.

Support e-mail address: <support@balabit.com>.

Support hotline: +36 1 398 6700 (available from 9 AM to 5 PM CET on weekdays)

B.5. Hardware specifications

SSB appliances are built on high performance, energy efficient, and reliable hardware that are easily mounted into standard rack mounts.

Product	Unit	Redundant Power Supply	Processor	Memory	Capacity	RAID	IPMI
SSB1000	1	No	INTEL X3430 2,4GHz QUAD	4 GB RAM	2x 1 TB SATA	Software raid	Yes
SSB1000d	2	Yes	2x INTEL XEON E5620 2,4GHz	6x DDR3 4GB	2x 1 TB SATA	Software raid	Yes
SSB5000	2	Yes	2x INTEL XEON E5620 2,4GHz	6x DDR3 4GB	12x 500 GB SATA	ADAPTEC 5405 SAS RAID Controller	Yes
SSB10000	2	Yes	2x INTEL XEON E5620 2,4GHz	6x DDR3 4GB	12x 1 TB SATA	ADAPTEC 5405 SAS RAID Controller	Yes

Table B.1. System related traps



Appendix C. syslog-ng Store Box Software Installation Guide

This leaflet describes how to install the syslog-ng Store Box (SSB) software on a certified hardware. The list of certified hardware is available at BalaBit.

C.1. Procedure – Installing the SSB software

Purpose:

To install a new SSB on a server, complete the following steps:

Steps:

- Step 1. Login to your [MyBalaBit account](#) and download the latest syslog-ng Store Box installation ISO file. Note that you need to have partner access to download syslog-ng Store Box ISO files. If you are a partner but do not see the ISO files, you can request partner access within MyBalaBit.
- Step 2. Mount the ISO image, or burn it to a CD-ROM.
- Step 3. Connect your computer to the *IPMI* interface of SSB. For details, see the following documents:
 - The *Onboard BMC/IPMI User's Guide*, available at [the BalaBit Hardware Documentation page](#).
- Step 4. Power on the server.
- Step 5. Login to the IPMI web interface, and boot the syslog-ng Store Box installation CD on the server using a virtual CD-ROM. For details, see the following documents:
 - The *Onboard BMC/IPMI User's Guide*, available at [the BalaBit Hardware Documentation page](#).
- Step 6. When the syslog-ng Store Box installer starts, select **Installer**, press Enter, and wait until the server finishes the boot process.
- Step 7. Select **Install a new SSB** and press **Enter** to start the installation process. Depending on the size of the disks, the installation process takes from a few minutes to an hour to complete. The progress of the installation is indicated in the **Installation Steps** window.
- Step 8. The installer displays the following question: **Warning, all data on the hard drive(s) will be erased. Are you sure?** Select **Yes** and press **Enter**.
- Step 9. The installer displays the MAC addresses of the network interfaces found in the SSB unit. Record these addresses.
- Step 10. If you are installing an SSB SAN Connect edition, the device includes an iSCSI HBA card. The installer will display the IQN number of the card during the installation. Record this number as it will be needed during the installation of the storage module.
- Step 11. The installer displays the product name, name of the SSB configuration that was installed, for example, *SSB1000*. If the product name displayed does not match the product you wanted to install, complete the following steps:



Step a. Check that the hardware configuration of the appliance matches the specifications provided by BalaBit.

Step b. If the configuration matches the specifications but the installer displays a different product name, contact the BalaBit Support Team.

Step 12. During the **Finishing the Setup** step, the installer performs RAID synchronization.

- Select **Yes** to perform the RAID synchronization. RAID synchronization is a two-step process, the progress of the active step is indicated on the progress bar. Wait until both steps are completed. Note that this synchronization takes several hours (about 8 hours on the average).
- Select **No** to skip the RAID synchronization. Note that the system will automatically perform the synchronization after the first boot, but in this case the process will take several days.

Step 13. After the installation is finished, press Enter to return to the main menu.

Step 14. Select **Reboot** and press Enter to restart the system. Wait until the system reboots.

Step 15. Connect your computer to the *EXT* interface of SSB. Create an alias IP address for your computer that falls into the *192.168.1.0/24* subnet (for example *192.168.1.10*). For details, see *Section 3.1, The initial connection to SSB (p. 17)*.

Step 16. Open the <http://192.168.1.1> URL in your web browser and verify that the Welcome Wizard of the syslog-ng Store Box is available.



Note

For details on the supported web browsers and operating systems, see *Section 4.1, Supported web browsers and operating systems (p. 33)*.

1. Welcome 2. License 3. Networking 4. Users 5. Certificate 6. Finish

Import old configuration

You can use your old exported configuration or continue clicking 'Next'.

Configuration: Browse... Upload

Encryption password:

Next

Figure C.1. The Welcome Wizard

Step 17. Power off the system.



Appendix D. syslog-ng Store Box VMware Installation Guide

This tutorial describes the possibilities and limitations of installing syslog-ng Store Box (SSB) 3 LTS as a virtual appliance under a VMware ESXi server.

D.1. Limitations of SSB under VMware

Version 3 LTS of SSB has no special support for running under VMware. While the basic functionality of SSB is not affected by running as a virtual appliance, the following limitations apply:

- SSB can only use fixed disk space assigned to the virtual host; it is not possible to use on-demand disk allocation scenarios. To increase the size of the virtual disk, see *Procedure D.3, Modifying the virtual disk size under VMware (p. 211)*
- SSB currently does not comply with the VMware Ready program and does not have vmware-tools installed. As a result, SSB does not support VMotion, nor any other VMware features that require vmware-tools. Consult the documentation of your VMware server to check which features work without vmware-tools.
- High availability mode is not supported in VMware.
- Hardware-related alerts and status indicators of SSB may display inaccurate information, for example, display degraded RAID status.

D.2. Procedure – Installing SSB under VMware ESXi

Purpose:

To install a new SSB under VMware ESXi, complete the following steps:



Warning

SSB can be installed under VMware ESXi 4.0 or later, earlier VMware versions are not supported.

Steps:

Step 1. Create the virtual machine for SSB using the following settings:

- Guest operating system: **Linux/Ubuntu 64bit**
- Allocate memory for the virtual machine. SSB requires a minimum of 512MB of memory, in addition to the memory limit of the indexed logspaces. The recommended size for the memory depends on the exact environment, but consider the following:
 - The base system requires 256MB
 - The syslog-ng server running on SSB requires about 128MB-1GB of memory, depending on the message load and on the configuration of SSB.



- For every log space, SSB requires additional memory to index the incoming messages. The amount of memory allocated for the indexer can be set individually for every log space.
- The hard disk controller must be **LSI Logic Parallel**.
- Do not use RAID for the hard disk, a single hard disk is sufficient for the system.
- SSB uses a single fixed-sized disk:
 - About 5GB is required for the base system, the remaining disk space is used to store data. Use one of the following disk sizes (in GB): 8, 12, 68, 160, 230, 250, 465, 500, 925, 1029, 2053, 3077, 4101, 4645, 5125, 6149, 7173, 8197, 9205.
 - It is not possible to use on-demand disk allocation scenarios. To increase the initial disk size, see *Procedure D.3, Modifying the virtual disk size under VMware* (p. 211)
- SSB requires 4 network cards, all of them must be **E1000**.

Step 2. After creating the virtual machine, edit the settings of the machine. Set the following options:

Step a. Under **Options > VMware Tools** enable the **Shutdown, Suspend, Reset** options, otherwise the SSB administrator will not be able to access these functions from the SSB web interface.

Step b. Under **Options > Boot options** enable the **Force BIOS Setup** option. This is required to be able to check the system time (and modify it if needed) before installing SSB.

Step 3. Login to your [MyBalaBit account](#) and download the latest syslog-ng Store Box installation ISO file. Note that you need to have purchased SSB as a virtual appliance or have partner access to download syslog-ng Store Box ISO files. If you are a partner but do not see the ISO files, you can request partner access within MyBalaBit.

Step 4. Mount the ISO image and boot the virtual machine. Follow the on-screen instructions to install SSB.

D.3. Procedure – Modifying the virtual disk size under VMware

SSB can only use fixed disk space assigned to the virtual host. If you must increase the size of the virtual disk, complete the following steps:

Step 1. Create a full system backup (configuration and data backup). For detailed instructions, see *Section 4.7, Data and configuration archiving and backups* (p. 54).

Step 2. Power down the virtual machine.

Step 3. Increase the storage size.

Step 4. Re-install SSB.

Step 5. Restore the system from the full backup. For detailed instructions, see *Procedure 14.7, Restoring SSB configuration and data* (p. 201).



Appendix E. License contract for BalaBit Product

SUBJECT OF THE LICENSE CONTRACT

This License Contract is entered into by and between BalaBit S.a.r.l. (or, based on your place of operation, one of its affiliates) as Licensor (hereinafter Company or Licensor, or BalaBit) and Licensee (hereinafter Licensee) and sets out the terms and conditions under which Licensee and/or Licensee's Authorized Subsidiaries may use the BalaBit product under this License Contract.

DEFINITIONS

In this License Contract, the following words shall have the following meanings:

Name	Description
Annexed Software	Any third party software that is not a BalaBit Product contained in the install media of the BalaBit Product.
BalaBit Product	Any software, hardware or service Licensed, sold, or provided by BalaBit including any installation, education, support and warranty services or any product falling under the copyright of BalaBit with the exception of the Annexed Software.
License Contract	The present BalaBit Product License Contract.
Product Documentation	Any documentation referring to the BalaBit Product or any module thereof, with special regard to the administration guide, the product description, the installation guide, user guides and manuals.
End-user Certificate	The document signed by Licensor which contains a) identification data of Licensee; b) configuration of BalaBit Product and designation of Licensed modules thereof; c) declaration of the parties on accepting the terms and conditions of this License Contract; and d) declaration of Licensee that is in receipt of the install media and the hardware appliance.
Product Usage Terms	Defines the conditions (related usage environment and limitations) under the BalaBit Product may be used by the Licensee.
Warranty Period	The period of twelve (12) months from the date of delivery of the BalaBit Product to Licensee.

Table E.1. Words and expressions

LICENSE GRANTS AND RESTRICTIONS

Based on the terms and conditions of the present License Contract and the applicable End-user Certificate, and the Product Usage Terms BalaBit grants to Licensee, a non-exclusive, perpetual license to use BalaBit Product. License is transferable only with the prior written approval of BalaBit.



Licensee shall use the BalaBit Product in accordance with the conditions sets by the Product Usage Terms, especially in the configuration and in the quantities specified in the End-user Certificate and Product Usage Terms.

On the install media (firmware CD-ROM, USB stick) all modules of the BalaBit Product will be presented, however, Licensee shall not be entitled to use any module which was not Licensed to it. Access rights to modules and IP connections are controlled by an "electronic key" accompanying the BalaBit Product.

Licensee shall be entitled to make one back-up copy of the install media containing the BalaBit Product.

Licensee shall make the BalaBit Product available solely to its own employees and those of the Authorized Subsidiaries.

Licensee shall take all reasonable steps to protect BalaBit's rights with respect to the BalaBit Product with special regard and care to protecting it from any unauthorized access.

Licensee shall, in 5 working days, properly answer the queries of BalaBit referring to the actual usage conditions of the BalaBit Product that may differ or allegedly differs from the License conditions.

Licensee shall not modify the BalaBit Product in any way, with special regard to the functions that inspect the usage of the software. Licensee shall install the code permitting the usage of the BalaBit Product according to the provisions defined for it by BalaBit. Licensee may not modify or cancel such codes. Configuration settings of the BalaBit Product in accordance with the possibilities offered by the system shall not be construed as modification of the software.

Licensee shall only be entitled to analyze the structure of the BalaBit Products (decompilation or reverse-engineering) if concurrent operation with a software developed by a third party is necessary, and upon request to supply the information required for concurrent operation BalaBit does not provide such information within 60 days from the receipt of such a request.

These user actions are limited to parts of the BalaBit Product which are necessary for concurrent operation. Any information obtained as a result of applying the previous Section (i) cannot be used for purposes other than concurrent operation with the BalaBit Product; (ii) cannot be disclosed to third parties unless it is necessary for concurrent operation with the BalaBit Product; (iii) cannot be used for the development, production or distribution of a different software which is similar to the BalaBit Product in its form of expression, or for any other act violating copyright.

For any Annexed Software contained by the same install media as the BalaBit Product, the terms and conditions defined by its copyright owner shall be properly applied. BalaBit does not grant any License rights to any Annexed Software.

Any usage of the BalaBit Product exceeding the limits and restrictions defined in this License Contract shall qualify as material breach of the License Contract and Licensee shall be fully liable towards BalaBit.

Licensee shall have the right to obtain and use content updates only if Licensee concludes a support contract that includes such content updates (maintenance of the software), or if Licensee has otherwise separately acquired the right to obtain and use such content updates. This License Contract does not otherwise permit Licensee to obtain and use content updates.



INTELLECTUAL PROPERTY RIGHTS

Licensee agrees that BalaBit owns all rights, titles, and interests related to the BalaBit Product, including all of BalaBit's patents, trademarks, trade names, inventions, economic intellectual property rights, know-how, and trade secrets relating to the design, manufacture, operation or service of the BalaBit Products.

The use by Licensee of any of these intellectual property rights is authorized only for the purposes set forth herein, and upon termination of this License Contract for any reason, such authorization shall cease.

The BalaBit Products are licensed only for the Licensee's own internal business purposes in every case, under the condition that such License does not convey any license, expressly or by implication, to manufacture duplicate or otherwise copy or reproduce any of the BalaBit Products. The sublicense to third parties, or provision of any services with the utilization of the BalaBit Product is not allowed. No other rights than expressly stated herein are granted to Licensee.

Licensee will take appropriate steps with its Authorized Subsidiaries, as BalaBit may request, to inform them of and assure compliance with the restrictions contained in the License Contract.

WARRANTIES

BalaBit warrants that during the Warranty Period, the magnetic or optical media upon which the BalaBit Product is recorded will not be defective under normal use. BalaBit will replace any defective media returned to it, accompanied by a dated proof of purchase, within the Warranty Period at no charge to Licensee. Upon receipt of the allegedly defective BalaBit Product, BalaBit will at its option, deliver a replacement BalaBit Product or BalaBit's current equivalent Product to Licensee at no additional cost. BalaBit will bear the delivery charges to Licensee for the replacement Product.

Should BalaBit Product be used in conjunction with third party software, BalaBit shall not be liable for errors due to third party software.

BalaBit warrants, that during the Warranty Period, the BalaBit Product, under normal use in the operating environment defined by BalaBit, and without unauthorized modification, will perform in substantial compliance with the Product Documentation accompanying the BalaBit Product, when used on that hardware for which it was installed, in compliance with the provisions of the user manuals and the recommendations of BalaBit.

The date of the notification sent to BalaBit shall qualify as the date of the failure. Licensee shall do its best to mitigate the consequences of that failure. If, during the Warranty Period, the BalaBit Product fails to comply with this warranty, and such failure is reported by Licensee to BalaBit within the Warranty Period, BalaBit's sole obligation and liability for breach of this warranty is, at BalaBit's sole option, either: (i) to correct such failure, (ii) to replace the defective BalaBit Product or (iii) to refund the license fees paid by Licensee for the applicable BalaBit Product.

BREACH OF CONTRACT

In case of breach of contract with respect to BalaBit, or the BalaBit Product, committed by violating any provision of the present License Contract, Licensee shall pay liquidated damages to BalaBit. The amount of the liquidated damages shall be twice as much as the price of the BalaBit Product concerned, on BalaBit's current Price List.



INTELLECTUAL PROPERTY INDEMNIFICATION

BalaBit shall pay all damages, costs and reasonable attorney's fees awarded against Licensee in connection with any claim brought against Licensee to the extent that such claim is based on a claim that Licensee's authorized use of the BalaBit Product infringes a patent, copyright, trademark or trade secret. Licensee shall notify BalaBit in writing of any such claim as soon as Licensee learns of it and shall cooperate fully with BalaBit in connection with the defense of that claim. BalaBit shall have sole control of that defense (including without limitation the right to settle the claim).

If Licensee is prohibited from using any BalaBit Product due to an infringement claim, or if BalaBit believes that any BalaBit Product is likely to become the subject of an infringement claim, BalaBit shall at its sole option, either: (i) obtain the right for Licensee to continue to use such BalaBit Product, (ii) replace or modify the BalaBit Product so as to make such BalaBit Product non-infringing and substantially comparable in functionality or (iii) refund to Licensee the amount paid for such infringing BalaBit Product and provide a pro-rated refund of any unused, prepaid maintenance fees paid by Licensee, in exchange for Licensee's return of such BalaBit Product to BalaBit.

Notwithstanding the above, BalaBit will have no liability for any infringement claim to the extent that it is based upon: (i) modification of the BalaBit Product other than by BalaBit, (ii) use of the BalaBit Product in combination with any product not specifically authorized by BalaBit to be combined with the BalaBit Product or (iii) use of the BalaBit Product in an unauthorized manner for which it was not designed.

LICENSE FEE

The End-user Certificate and the Product Usage Term contain the details of the purchased License and usage limitations. This information serves as the calculation base of the License fee.

Licensee acknowledges that payment of the License fees is a condition of lawful usage.

License fees do not contain any installation or post charges, taxes, duties, etc.

The license right of BalaBit Product is transferred to the Licensee only when Licensee pays the License fee to BalaBit. In case of non-payment BalaBit has right to terminate, or rescind from the License Contract with immediate effect and Licensee has to send back the BalaBit Product on their own cost and takes all liability regarding the unlawful usage and the early termination.

DISCLAIMER OF WARRANTIES

EXCEPT AS SET OUT IN THIS LICENSE CONTRACT, BALABIT MAKES NO WARRANTIES OF ANY KIND WITH RESPECT TO THE BALABIT PRODUCT. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW, BALABIT EXCLUDES ANY OTHER WARRANTIES, INCLUDING BUT NOT LIMITED TO ANY IMPLIED WARRANTIES OF SATISFACTORY QUALITY, MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS.

LIMITATION OF LIABILITY

SOME STATES AND COUNTRIES, INCLUDING MEMBER COUNTRIES OF THE EUROPEAN UNION, DO NOT ALLOW THE LIMITATION OR EXCLUSION OF LIABILITY FOR INCIDENTAL OR CONSEQUENTIAL DAMAGES AND, THEREFORE, THE FOLLOWING LIMITATION OR EXCLUSION MAY NOT APPLY TO THIS LICENSE CONTRACT IN THOSE STATES AND COUNTRIES. TO THE MAXIMUM EXTENT PERMITTED BY APPLICABLE LAW AND REGARDLESS OF WHETHER ANY



REMEDY SET OUT IN THIS LICENSE CONTRACT FAILS OF ITS ESSENTIAL PURPOSE, IN NO EVENT SHALL BALABIT BE LIABLE TO LICENSEE FOR ANY SPECIAL, CONSEQUENTIAL, INDIRECT OR SIMILAR DAMAGES OR LOST PROFITS OR LOST DATA ARISING OUT OF THE USE OR INABILITY TO USE THE BALABIT PRODUCT EVEN IF BALABIT HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

IN NO CASE SHALL BALABIT'S TOTAL LIABILITY UNDER THIS LICENSE CONTRACT EXCEED THE FEES PAID BY LICENSEE FOR THE BALABIT PRODUCT LICENSED UNDER THIS LICENSE CONTRACT.

BalaBit shall not take any responsibility, for damages caused by the usage of the BalaBit Product which is not in accordance with the Product Usage Terms.

DURATION AND TERMINATION

This License Contract shall come into effect on the day when the End-user Certificate and the declaration of the Licensee on accepting the terms and conditions of this License Contract and the declaration of Licensee that is in receipt of the install media and the hardware appliance are both signed by the duly authorized representatives of the relevant party.

Licensee may terminate the License Contract at any time by written notice sent to BalaBit and by simultaneously destroying all copies of the BalaBit Product licensed under this License Contract.

BalaBit may terminate this License Contract with immediate effect by written notice to Licensee, if Licensee is in material or persistent breach of the License Contract and either that breach is incapable of remedy or Licensee shall have failed to remedy that breach within 30 days after receiving written notice requiring it to remedy that breach.

AMENDMENTS

Save as expressly provided in this License Contract, no amendment or variation of this License Contract shall be effective unless in writing and signed by a duly authorized representative of the parties to it.

WAIVER

The failure of a party to exercise or enforce any right under this License Contract shall not be deemed to be a waiver of that right nor operate to bar the exercise or enforcement of it at any time or times thereafter.

SEVERABILITY

If any part of this License Contract becomes invalid, illegal or unenforceable, the parties shall in such an event negotiate in good faith in order to agree on the terms of a mutually satisfactory provision to be substituted for the invalid, illegal or unenforceable provision which as nearly as possible validly gives effect to their intentions as expressed in this License Contract.

NOTICES

Any notice required to be given pursuant to this License Contract shall be in writing and shall be given by delivering the notice by hand, or by sending the same by prepaid first class post (airmail if to an address outside the country of posting) or via e-mail to the address of the relevant party. Any notice given according to the above procedure shall be deemed to have been given at the time of delivery (if delivered by hand) and when received (if sent by post or via e-mail).



APPLICABLE LAW AND LANGUAGE AND SETTLEMENT OF LEGAL DISPUTES

This agreement shall be governed by and construed in accordance with the laws of the Hungary without regard for its conflicts of law provisions. The language for all communications regarding this Agreement shall be English. The Parties agree that any dispute arising from this Agreement, or the breach, termination, validity or interpretation thereof or in relation thereto shall come under the exclusive jurisdiction of the Hungarian Court as defined below.

The Parties irrevocably agree that any dispute, controversy or claim arising out of or in connection with this agreement, or the breach, termination or invalidity thereof, shall be finally settled by arbitration by the Permanent Arbitration Court attached to the Hungarian Chamber of Commerce and Industry acting in accordance with its own Rules of Procedure. The place of arbitration shall be Budapest, the number of arbitrators shall be three (3) and the language to be used in the arbitral proceedings shall be English.

The Arbitration Court shall consist of three arbitrators, out of which one shall act as chairman. The chairman should have the competence of judgeship. The Arbitration Court shall be created in the manner that the prosecuting Party, indicating the subject of the debate and nominating an arbitrator, calls the counterparty in writing to nominate the other arbitrator and the nominated arbitrators will elect the chairman.

MISCELLANEOUS

Headings are for convenience only and shall be ignored in interpreting this License Contract.

This License Contract and the rights granted in this License Contract may not be assigned, sublicensed or otherwise transferred in whole or in part by Licensee without BalaBit's prior written consent.

An independent third party auditor, reasonably acceptable to BalaBit and Licensee, may upon reasonable notice to Licensee and during normal business hours, but not more often than once each year, inspect Licensee's relevant records in order to confirm that usage of the BalaBit Product complies with the terms and conditions of this License Contract. BalaBit shall bear the costs of such audit. All audits shall be subject to the reasonable safety and security policies and procedures of Licensee.

In case of non-acceptance in the person, an auditor, appointed by BalaBit shall keep full, complete and accurate inspect concerning that usage of the BalaBit Product complies with the terms and conditions of this License Contract.

The auditor shall be entitled to examine, inspect, copy and audit the usage of the BalaBit Product. If the inspection or audit reveals that the usage is not complies with the conditions of the License Contract the Licensee shall immediately:

- (a) pay to BalaBit the amount of any underpayment, together with interest on that amount calculated at the rate of two per cent (2%) over the Barclay Bank base rate from time to time; and
- (b) pay the costs of the audit and/or inspection where that audit or inspection reveals an underpayment in excess of five per cent (5%).

In case of the License shall not let the auditor to inspect, or examine the usage of BalaBit Product, BalaBit has right to terminate or rescind from the License Contract with immediate effect and Licensee has to send back the BalaBit Product on their own cost and takes all liability regarding the unlawful usage and the early termination.

This License Contract constitutes the entire agreement between the parties with regard to the subject matter hereof.



Appendix F. Creative Commons Attribution Non-commercial No Derivatives (by-nc-nd) License

THE WORK (AS DEFINED BELOW) IS PROVIDED UNDER THE TERMS OF THIS CREATIVE COMMONS PUBLIC LICENSE ("CCPL" OR "LICENSE"). THE WORK IS PROTECTED BY COPYRIGHT AND/OR OTHER APPLICABLE LAW. ANY USE OF THE WORK OTHER THAN AS AUTHORIZED UNDER THIS LICENSE OR COPYRIGHT LAW IS PROHIBITED. BY EXERCISING ANY RIGHTS TO THE WORK PROVIDED HERE, YOU ACCEPT AND AGREE TO BE BOUND BY THE TERMS OF THIS LICENSE. TO THE EXTENT THIS LICENSE MAY BE CONSIDERED TO BE A CONTRACT, THE LICENSOR GRANTS YOU THE RIGHTS CONTAINED HERE IN CONSIDERATION OF YOUR ACCEPTANCE OF SUCH TERMS AND CONDITIONS.

1. Definitions

- a. "Adaptation" means a work based upon the Work, or upon the Work and other pre-existing works, such as a translation, adaptation, derivative work, arrangement of music or other alterations of a literary or artistic work, or phonogram or performance and includes cinematographic adaptations or any other form in which the Work may be recast, transformed, or adapted including in any form recognizably derived from the original, except that a work that constitutes a Collection will not be considered an Adaptation for the purpose of this License. For the avoidance of doubt, where the Work is a musical work, performance or phonogram, the synchronization of the Work in timed-relation with a moving image ("synching") will be considered an Adaptation for the purpose of this License.
- b. "Collection" means a collection of literary or artistic works, such as encyclopedias and anthologies, or performances, phonograms or broadcasts, or other works or subject matter other than works listed in Section 1(f) below, which, by reason of the selection and arrangement of their contents, constitute intellectual creations, in which the Work is included in its entirety in unmodified form along with one or more other contributions, each constituting separate and independent works in themselves, which together are assembled into a collective whole. A work that constitutes a Collection will not be considered an Adaptation (as defined above) for the purposes of this License.
- c. "Distribute" means to make available to the public the original and copies of the Work through sale or other transfer of ownership.
- d. "Licensor" means the individual, individuals, entity or entities that offer(s) the Work under the terms of this License.
- e. "Original Author" means, in the case of a literary or artistic work, the individual, individuals, entity or entities who created the Work or if no individual or entity can be identified, the publisher; and in addition (i) in the case of a performance the actors, singers, musicians, dancers, and other persons who act, sing, deliver, declaim, play in, interpret or otherwise perform literary or artistic works or expressions of folklore; (ii) in the case of a phonogram the producer being the person or legal entity who first fixes the sounds of a performance or other sounds; and, (iii) in the case of broadcasts, the organization that transmits the broadcast.



- f. "Work" means the literary and/or artistic work offered under the terms of this License including without limitation any production in the literary, scientific and artistic domain, whatever may be the mode or form of its expression including digital form, such as a book, pamphlet and other writing; a lecture, address, sermon or other work of the same nature; a dramatic or dramatico-musical work; a choreographic work or entertainment in dumb show; a musical composition with or without words; a cinematographic work to which are assimilated works expressed by a process analogous to cinematography; a work of drawing, painting, architecture, sculpture, engraving or lithography; a photographic work to which are assimilated works expressed by a process analogous to photography; a work of applied art; an illustration, map, plan, sketch or three-dimensional work relative to geography, topography, architecture or science; a performance; a broadcast; a phonogram; a compilation of data to the extent it is protected as a copyrightable work; or a work performed by a variety or circus performer to the extent it is not otherwise considered a literary or artistic work.
 - g. "You" means an individual or entity exercising rights under this License who has not previously violated the terms of this License with respect to the Work, or who has received express permission from the Licensor to exercise rights under this License despite a previous violation.
 - h. "Publicly Perform" means to perform public recitations of the Work and to communicate to the public those public recitations, by any means or process, including by wire or wireless means or public digital performances; to make available to the public Works in such a way that members of the public may access these Works from a place and at a place individually chosen by them; to perform the Work to the public by any means or process and the communication to the public of the performances of the Work, including by public digital performance; to broadcast and rebroadcast the Work by any means including signs, sounds or images.
 - i. "Reproduce" means to make copies of the Work by any means including without limitation by sound or visual recordings and the right of fixation and reproducing fixations of the Work, including storage of a protected performance or phonogram in digital form or other electronic medium.
2. *Fair Dealing Rights.* Nothing in this License is intended to reduce, limit, or restrict any uses free from copyright or rights arising from limitations or exceptions that are provided for in connection with the copyright protection under copyright law or other applicable laws.
3. *License Grant.* Subject to the terms and conditions of this License, Licensor hereby grants You a worldwide, royalty-free, non-exclusive, perpetual (for the duration of the applicable copyright) license to exercise the rights in the Work as stated below:
- a. to Reproduce the Work, to incorporate the Work into one or more Collections, and to Reproduce the Work as incorporated in the Collections; and,
 - b. to Distribute and Publicly Perform the Work including as incorporated in Collections.
- The above rights may be exercised in all media and formats whether now known or hereafter devised. The above rights include the right to make such modifications as are technically necessary to exercise the rights in other media and formats, but otherwise you have no rights to make Adaptations. Subject to 8(f), all rights not expressly granted by Licensor are hereby reserved, including but not limited to the rights set forth in Section 4(d).
4. *Restrictions.* The license granted in Section 3 above is expressly made subject to and limited by the following restrictions:



- a. You may Distribute or Publicly Perform the Work only under the terms of this License. You must include a copy of, or the Uniform Resource Identifier (URI) for, this License with every copy of the Work You Distribute or Publicly Perform. You may not offer or impose any terms on the Work that restrict the terms of this License or the ability of the recipient of the Work to exercise the rights granted to that recipient under the terms of the License. You may not sublicense the Work. You must keep intact all notices that refer to this License and to the disclaimer of warranties with every copy of the Work You Distribute or Publicly Perform. When You Distribute or Publicly Perform the Work, You may not impose any effective technological measures on the Work that restrict the ability of a recipient of the Work from You to exercise the rights granted to that recipient under the terms of the License. This Section 4(a) applies to the Work as incorporated in a Collection, but this does not require the Collection apart from the Work itself to be made subject to the terms of this License. If You create a Collection, upon notice from any Licensor You must, to the extent practicable, remove from the Collection any credit as required by Section 4(c), as requested.
- b. You may not exercise any of the rights granted to You in Section 3 above in any manner that is primarily intended for or directed toward commercial advantage or private monetary compensation. The exchange of the Work for other copyrighted works by means of digital file-sharing or otherwise shall not be considered to be intended for or directed toward commercial advantage or private monetary compensation, provided there is no payment of any monetary compensation in connection with the exchange of copyrighted works.
- c. If You Distribute, or Publicly Perform the Work or Collections, You must, unless a request has been made pursuant to Section 4(a), keep intact all copyright notices for the Work and provide, reasonable to the medium or means You are utilizing: (i) the name of the Original Author (or pseudonym, if applicable) if supplied, and/or if the Original Author and/or Licensor designate another party or parties (for example a sponsor institute, publishing entity, journal) for attribution ("Attribution Parties") in Licensor's copyright notice, terms of service or by other reasonable means, the name of such party or parties; (ii) the title of the Work if supplied; (iii) to the extent reasonably practicable, the URI, if any, that Licensor specifies to be associated with the Work, unless such URI does not refer to the copyright notice or licensing information for the Work. The credit required by this Section 4(c) may be implemented in any reasonable manner; provided, however, that in the case of a Collection, at a minimum such credit will appear, if a credit for all contributing authors of Collection appears, then as part of these credits and in a manner at least as prominent as the credits for the other contributing authors. For the avoidance of doubt, You may only use the credit required by this Section for the purpose of attribution in the manner set out above and, by exercising Your rights under this License, You may not implicitly or explicitly assert or imply any connection with, sponsorship or endorsement by the Original Author, Licensor and/or Attribution Parties, as appropriate, of You or Your use of the Work, without the separate, express prior written permission of the Original Author, Licensor and/or Attribution Parties.
- d. For the avoidance of doubt:
 - i. Non-waivable Compulsory License Schemes. In those jurisdictions in which the right to collect royalties through any statutory or compulsory licensing scheme cannot be waived, the Licensor reserves the exclusive right to collect such royalties for any exercise by You of the rights granted under this License;
 - ii. Waivable Compulsory License Schemes. In those jurisdictions in which the right to collect royalties through any statutory or compulsory licensing scheme can be waived, the Licensor reserves the exclusive right to collect such royalties for any exercise by You of the rights



granted under this License if Your exercise of such rights is for a purpose or use which is otherwise than noncommercial as permitted under Section 4(b) and otherwise waives the right to collect royalties through any statutory or compulsory licensing scheme; and,

iii. Voluntary License Schemes. The Licensor reserves the right to collect royalties, whether individually or, in the event that the Licensor is a member of a collecting society that administers voluntary licensing schemes, via that society, from any exercise by You of the rights granted under this License that is for a purpose or use which is otherwise than noncommercial as permitted under Section 4(b).

e. Except as otherwise agreed in writing by the Licensor or as may be otherwise permitted by applicable law, if You Reproduce, Distribute or Publicly Perform the Work either by itself or as part of any Collections, You must not distort, mutilate, modify or take other derogatory action in relation to the Work which would be prejudicial to the Original Author's honor or reputation.

5. *Representations, Warranties and Disclaimer* UNLESS OTHERWISE MUTUALLY AGREED BY THE PARTIES IN WRITING, LICENSOR OFFERS THE WORK AS-IS AND MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND CONCERNING THE WORK, EXPRESS, IMPLIED, STATUTORY OR OTHERWISE, INCLUDING, WITHOUT LIMITATION, WARRANTIES OF TITLE, MERCHANTIBILITY, FITNESS FOR A PARTICULAR PURPOSE, NONINFRINGEMENT, OR THE ABSENCE OF LATENT OR OTHER DEFECTS, ACCURACY, OR THE PRESENCE OF ABSENCE OF ERRORS, WHETHER OR NOT DISCOVERABLE. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OF IMPLIED WARRANTIES, SO SUCH EXCLUSION MAY NOT APPLY TO YOU.

6. *Limitation on Liability*. EXCEPT TO THE EXTENT REQUIRED BY APPLICABLE LAW, IN NO EVENT WILL LICENSOR BE LIABLE TO YOU ON ANY LEGAL THEORY FOR ANY SPECIAL, INCIDENTAL, CONSEQUENTIAL, PUNITIVE OR EXEMPLARY DAMAGES ARISING OUT OF THIS LICENSE OR THE USE OF THE WORK, EVEN IF LICENSOR HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

7. *Termination*

a. This License and the rights granted hereunder will terminate automatically upon any breach by You of the terms of this License. Individuals or entities who have received Collections from You under this License, however, will not have their licenses terminated provided such individuals or entities remain in full compliance with those licenses. Sections 1, 2, 5, 6, 7, and 8 will survive any termination of this License.

b. Subject to the above terms and conditions, the license granted here is perpetual (for the duration of the applicable copyright in the Work). Notwithstanding the above, Licensor reserves the right to release the Work under different license terms or to stop distributing the Work at any time; provided, however that any such election will not serve to withdraw this License (or any other license that has been, or is required to be, granted under the terms of this License), and this License will continue in full force and effect unless terminated as stated above.

8. *Miscellaneous*

a. Each time You Distribute or Publicly Perform the Work or a Collection, the Licensor offers to the recipient a license to the Work on the same terms and conditions as the license granted to You under this License.

b. If any provision of this License is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this License, and without further



action by the parties to this agreement, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

- c. No term or provision of this License shall be deemed waived and no breach consented to unless such waiver or consent shall be in writing and signed by the party to be charged with such waiver or consent.
- d. This License constitutes the entire agreement between the parties with respect to the Work licensed here. There are no understandings, agreements or representations with respect to the Work not specified here. Licensor shall not be bound by any additional provisions that may appear in any communication from You. This License may not be modified without the mutual written agreement of the Licensor and You.
- e. The rights granted under, and the subject matter referenced, in this License were drafted utilizing the terminology of the Berne Convention for the Protection of Literary and Artistic Works (as amended on September 28, 1979), the Rome Convention of 1961, the WIPO Copyright Treaty of 1996, the WIPO Performances and Phonograms Treaty of 1996 and the Universal Copyright Convention (as revised on July 24, 1971). These rights and subject matter take effect in the relevant jurisdiction in which the License terms are sought to be enforced according to the corresponding provisions of the implementation of those treaty provisions in the applicable national law. If the standard suite of rights granted under applicable copyright law includes additional rights not granted under this License, such additional rights are deemed to be included in the License; this License is not intended to restrict the license of any rights under applicable law.



Glossary

alias IP	An additional IP address assigned to an interface that already has an IP address. The normal and alias IP addresses both refer to the same physical interface.
auditing policy	The auditing policy determines which events are logged on host running Microsoft Windows operating systems.
authentication	The process of verifying the authenticity of a user or client before allowing access to a network system or service.
BSD-syslog protocol	The old syslog protocol standard described in RFC 3164 <i>The BSD syslog Protocol</i> . Sometimes also referred to as the legacy-syslog protocol.
CA	A Certificate Authority (CA) is an institute that issues certificates.
certificate	A certificate is a file that uniquely identifies its owner. Certificates contains information identifying the owner of the certificate, a public key itself, the expiration date of the certificate, the name of the CA that signed the certificate, and some other data.
client mode	In client mode, syslog-ng collects the local logs generated by the host and forwards them through a network connection to the central syslog-ng server or to a relay.
destination	A logspace or a remote database or server where the log messages are stored.
destination driver	A communication method that syslog-ng uses to send log messages to a destination, for example to a remote server or to the hard disk.
destination, remote	A destination that sends log messages to a remote host (that is, a syslog-ng relay or server) using a network connection.
destination, local	A destination that transfers log messages to a logspace.
disk buffer	The Premium Edition of syslog-ng can store messages on the local hard disk if the central log server or the network connection to the server becomes unavailable.
disk queue	See <i>disk buffer</i> .
domain name	The name of a network, for example <i>balabit.com</i> .
External network interface	The <i>external</i> interface (labeled <i>EXT</i>) is used for general communication between the clients and the servers. If the management interface is not configured, the external interface is used for management purposes as well.



filter	An expression that selects only those message from a source that match the conditions set in the filter.
firmware	A firmware is a collection of the software components running on SSB. Individual software components cannot be upgraded on SSB, only the entire firmware. SSB contains two firmwares, an external (or boot) firmware and an internal (or core) firmware. These can be upgraded separately.
gateway	A device that connect two or more parts of the network, for example your local intranet and the external network (the Internet). Gateways act as entrances into other networks.
High Availability	High Availability (HA) uses a second SSB unit (called slave node) to ensure that the services are available even if the first unit (called master node) breaks down.
host	A computer connected to the network.
hostname	A name that identifies a host on the network. Hostnames can contain only alphanumerical characters (A-Z, a-z, 0-9) and the hyphen (-) character.
HA network interface	The <i>HA</i> interface (labeled <i>HA</i>) is an interface reserved for communication between the nodes of SSB clusters.
IETF-syslog protocol	The syslog-protocol standard developed by the Internet Engineering Task Force (IETF), described in RFC 5424 <i>The IETF syslog Protocol</i> .
key pair	A private key and its related public key. The private key is known only to the owner; the public key can be freely distributed. Information encrypted with the private key can only be decrypted using the public key.
LDAP	The Lightweight Directory Access Protocol (LDAP), is an application protocol for querying and modifying data using directory services running over TCP/IP.
log path	A combination of sources, filters, parsers, rewrite rules, and destinations: syslog-ng examines all messages arriving to the sources of the logpath and sends the messages matching all filters to the defined destinations.
logstore	A binary logfile format that can encrypt, sign, compress, and timestamp log messages.
log source host	A host or network device (including syslog-ng clients and relays) that sends logs to the syslog-ng Store Box. Log source hosts can be servers, routers, desktop computers, or other devices capable of sending syslog messages or running syslog-ng.
LSH	See <i>log source host</i> .



Management interface	network	The <i>management</i> interface (labeled <i>MGMT</i>) is used exclusively for communication between SSB and the auditor or the administrator of the syslog-ng Store Box.
master node		The active SSB unit that is collecting the log messages when SSB is used in High Availability mode.
name server		A network computer storing the IP addresses corresponding to domain names.
node		An SSB unit running in High Availability mode.
output buffer		A part of the memory of the host where syslog-ng stores outgoing log messages if the destination cannot accept the messages immediately.
output queue		Messages from the output queue are sent to the target syslog-ng server. The syslog-ng application puts the outgoing messages directly into the output queue, unless the output queue is full. The output queue can hold 64 messages, this is a fixed value and cannot be modified.
overflow queue		See <i>output buffer</i> .
ping		A command that sends a message from a host to another host over a network to test connectivity and packet loss.
port		A number ranging from 1 to 65535 that identifies the destination application of the transmitted data. For example: SSH commonly uses port 22, web servers (HTTP) use port 80, and so on.
Public-key authentication		An authentication method that uses encryption key pairs to verify the identity of a user or a client.
redundant Heartbeat interface		A redundant Heartbeat interface is a virtual interface that uses an existing interface of the SSB device to detect that the other node of the SSB cluster is still available. The virtual interface is not used to synchronize data between the nodes, only Heartbeat messages are transferred.
regular expression		A regular expression is a string that describes or matches a set of strings. The syslog-ng application supports extended regular expressions (also called POSIX modern regular expressions).
relay mode		In relay mode, syslog-ng receives logs through the network from syslog-ng clients and forwards them to the central syslog-ng server using a network connection.
SSB		An abbreviation of the syslog-ng Store Box name.
slave node		The passive SSB unit that replaces the active unit (the master node) if the master becomes unavailable.
source		A way for SSB to receive syslog messages.



source, network	A source that receives log messages from a remote host using a network connection. The UDP, TCP, and TLS methods are supported using the BSD-syslog and the IETF-syslog protocols.
source, local	A source that receives log messages locally from SSB.
source driver	A communication method used to receive log messages.
SNMP	Simple Network Management Protocol (SNMP) is an industry standard protocol used for network management. SSB can receive SNMP messages from remote hosts and convert them to syslog messages, and can also send its own SNMP traps to a central SNMP server.
split brain	A split brain situation occurs when for some reason (for example the loss of connection between the nodes) both nodes of a SSB cluster become active (master). This might cause that new data (for example log messages) is created on both nodes without being replicated to the other node. Thus, it is likely in this situation that two diverging sets of data are created, which cannot be trivially merged.
syslog-ng	The syslog-ng application is a flexible and highly scalable system logging application, typically used to manage log messages and implement centralized logging.
syslog-ng agent	The syslog-ng agent for Windows is a log collector and forwarder application for the Microsoft Windows platform. It collects the log messages of the Windows-based host and forwards them to SSB using regular or SSL-encrypted TCP connections.
syslog-ng client	A host running syslog-ng in client mode.
syslog-ng Premium Edition	The syslog-ng Premium Edition is the commercial version of the open-source application. It offers additional features, like encrypted message transfer and an agent for Microsoft Windows platforms.
syslog-ng relay	A host running syslog-ng in relay mode.
syslog-ng server	A host running syslog-ng in server mode, like SSB.
TLS	Transport Layer Security (TLS) and its predecessor, Secure Sockets Layer (SSL), are cryptographic protocols which provide secure communications on the Internet. The syslog-ng application can encrypt the communication between the clients and the server using TLS to prevent unauthorized access to sensitive log messages.
template	A user-defined structure that can be used to restructure log messages or automatically generate file names.
traceroute	A command that shows all routing steps (the path of a message) between two hosts.



Index

unmapping SAN volumes, 113

A

- accessing SSB using SSH, 101
- accounting SSB, 83, 159
- admin password (see administrator password)
- administrator password, 26
- alerts
 - master, 51
 - message rate, 50, 118
- alias interfaces, 38
- alias IP addresses, 18, 21, 38
- archiving, 97
 - log messages, 59
 - NetApp devices, 56, 60, 65
 - notifications, 176
 - Windows 2008 R2, 56, 60, 65
- artificial ignorance, 183
 - message classification, 188
- auditing configuration changes, 83
- authentication, 7-8
 - LDAP, 69
 - to Microsoft Active Directory, 69, 73
 - to RADIUS servers, 76
 - users, 69, 73

B

- backups, 54
 - encrypting, 62
 - file ownerships, 68
 - file permissions, 68
 - NetApp devices, 56, 60, 65
 - notifications, 176
 - restore, 201
 - Windows 2008 R2, 56, 60, 65
- browser requirements, 33
- browsers, 33
 - supported versions, 33
- browsing log messages, 156
- browsing reports, 178
- built-in sources, 114

C

- certificates, 7
 - accepted formats, 108
 - changing, 105
 - extendedKeyUsage, 108
 - for TLS authentication, 153
 - LDAP servers, 75
 - managing, 105
 - Timestamping Authority, 105
 - uploading, 108
 - web interface, 27, 105
 - X509v3 Extended Key Usage, 108
- changelogs, 83, 159
- changing certificates
 - Timestamping Authority, 105
- classifying messages, 183
 - alerts, 176
 - pattern matching concepts, 185
- collecting debug information, 194
- collecting system-state information, 194
- commit log, 83
- compliance, 2
- configuration
 - backups, 54
 - changes, 83
 - delimiters, 123
 - log paths, 144
 - logspaces, 121
 - network interfaces, 37
 - SAN access, 31
 - sources, 116
- console menu, 100
- controlling SSB
 - rebooting, 84
 - shutting down, 84
- converting SNMP to syslog, 114
- core files, 193
- creating local destinations, 121
- creating log spaces on volumes, 119, 121
- creating logspaces, 121
- creating sources, 116
- custom log files, 121
- custom reports, 180
- custom sources, 116

D

- dashboard, 177, 196



- database format, 138
- database templates, 138
- date and time, 41
 - configuring, 42
- debug logging, 194
- default reports, 180
- default sources, 114
- default usergroups, 81
- deleting
 - log files, 60, 127
 - search filters, 166
- destinations, 3, 119, 135
- disabling message parsing, 118
- displaying selected messages, 159
- DNS
 - server, 40
- downgrading the firmware
 - rollback, 95
- DRBD
 - adjusting synchronization speed, 88
- DRBD status, 87
 - connected, 87
 - invalidated, 87
 - split brain, 88, 198
 - sync source, 87
 - sync target, 87
 - wfconnection, 88

E

- e-mail alerts, 42-43, 47-48
- e-mailing reports, 43
- encrypting log messages, 7
- exporting
 - search results, 159
 - SSB configuration, 97
- exporting pattern database, 188
- external timestamps, 151

F

- facilities, 12, 14
- feature releases, 9
- file destinations, 119
- filtering messages, 146
- filtering search results, 159
- filters, 4, 146
 - available for every user, 167
 - global, 167
 - predefined, 166

- finding patterns, 186
- firmware, 9
 - high availability, 9
 - rollback, 95
 - update, 93, 95
- flow-control, 5
 - multiple destinations, 7

G

- GPG, 62
- group management
 - local, 72

H

- HA (see High Availability)
- hardware serial number, 105
- High Availability, 8
 - address, 26, 40
 - adjusting synchronization speed, 88
 - installation, 205
 - log messages, 85
 - manual takeover, 85
 - next-hop monitoring, 91
 - Node HA status, 86
 - Node HA UUID, 86
 - node replacement, 200
 - reboot cluster, 86
 - recovery, 197
 - redundant Heartbeat interfaces, 89-90
 - status, 86
 - synchronizing time, 42
- high availability mode, 84
- history of changes, 159
- hostlists, 110, 115, 129, 132
 - changing, 111-112
 - creating new, 110
 - importing, 111
 - modifying, 111-112

I

- ILOM, 104
- importing
 - certificates, 108
 - SSB configuration, 98
- importing certificates, 153
- importing pattern database, 188
- indexing



- delimiters, 123
- installing
 - SSB, 204
- Integrated Lights Out Management, 104
- Intelligent Platform Management Interface, 104
- IPMI, 104

J

- JavaScript, 33

L

- LDAP authentication, 69, 73
- LDAP servers
 - certificates, 75
 - failover, 74
 - Microsoft Active Directory on Windows 2000, 75
 - mutual authentication, 75
 - Windows 2000, 75
- license, 10
 - update, 96
- local console, 100
- local name resolution, 152
- local SSB users, 69
- local time, 13, 15
- local users
 - password management, 70
 - usergroups, 72
- lock management, 37
- log
 - debug mode, 194
 - system state, 194
 - tailing, 193
 - viewing, 193
- log message structure
 - BSD-syslog protocol, 11
 - IETF-syslog protocol, 13
 - legacy-syslog protocol, 11
 - RFC 3164, 11
 - RFC 5424, 13
- log messages
 - alerts, 176
 - browsing, 156
 - exporting search results, 159
 - filtering search results, 159
 - reports, 178
 - searching, 156
 - structure of, 11
- log paths, 3

- creating new, 144
- defaults, 144
- flow-control, 5
- log space size, 127
- log spaces
 - accessing shares, 133
 - deleting log files, 127
 - instant archiving, 127
 - instant backup, 127
 - logstore limitations, 120
 - manual backup, 127
 - sharing, 128
 - size of, 127
 - using logcat, 121
- log statements (see log paths)
- logcat, 121
- logging procedure, 4
- logspace, 119
- logstore, 119, 121
 - browsing encrypted, 171
 - cipher method used for encrypting, 151
 - digest method used for encrypting, 151
 - limitations of, 120
 - timestamping authority, 151
 - using logcat, 121
- Long Term Supported releases, 9
- LTS releases, 9

M

- Management Information Base, 47
- managing certificates
 - Timestamping Authority, 105
- managing search filters, 166
- message classification, 183
 - alerts, 176
- message destinations, 119, 135
- Message dialog, 37
- message facilities, 12, 14
- message filtering
 - using parsers, 190
- message filters, 146
- message rate alerting, 50, 118
- message sources, 114
- MIB, 47
- Microsoft Active Directory
 - supported platforms, 75
- Microsoft Windows Network, 63
- monitoring, 43-44, 47, 48



- MIB, 47
- mounting shares, 133
- mutual authentication, 8

N

- name resolution, 152
 - local, 152
- NetApp, 56, 60, 65
- network interfaces
 - alias interfaces, 38
 - alias IP addresses, 38
 - configuring, 37
 - configuring interface speed, 39
 - external interface, 8, 17, 38
 - HA interface, 8
 - IPMI interface, 8
 - management interface, 8, 39-40
- network shares, 128
- next-hop router monitoring, 91
- NFS, 128, 133
- no-parse, 118
- NTP servers, 42
- number of active hosts, 35
- number of active senders, 35

O

- operational-report, 180
- out-of-band management, 104
- output buffer, 6

P

- parameters
 - log_fetch_limit() , 5
 - log_fifo_size() , 5
 - log_iw_size() , 6
 - max_connections() , 6
- parsing messages, 188
 - filtering parsed messages, 190
- password
 - admin, 26
 - changing the root, 102
 - root, 26, 101
- password policies, 70
- pattern database, 183
 - adding patterns, 186
 - browsing, 186
 - create ruleset, 186

- creating parsers, 188
- export database, 188
- export ruleset, 188
- import database, 188
- import ruleset, 188
- structure of, 183
- using the results, 190
- pattern databases
 - pattern matching precedence, 185
- pattern matching, 183
 - procedure of, 185
 - searching for patterns, 186
- patternldb (see pattern database)
- persistent name resolution, 152
- ping, 192
- predefined filters, 166
- preferences, 34
- preventing message loss (see flow-control)
- private key
 - accepted formats, 29, 108
- public-key authentication on SSB, 101

Q

- querying SSB via SNMP, 46

R

- RADIUS authentication, 76
- reboot, 84
- redundant Heartbeat interfaces, 89-90
- redundant Heartbeat status
 - degraded, 91
 - degraded-working, 91
 - invalid, 91
 - not used, 91
 - ok, 91
- releases, 9
- remote destinations, 135
 - remote servers, 139
 - SNMP, 142
 - SQL databases, 135
- remote server, 135
- reports, 178, 180
 - browsing, 178
 - contents, 180
 - custom reports, 180
 - default, 180
 - e-mailing, 43
- restart, 84



- restoring a backup, **201**
- results of search (see search results)
- reverting the firmware version, 95
- root password, 26, 101-102
- Routing table, 40
 - management traffic, 40
- rsync over SSH, 63
- S**
- samba shares, 128, 133
- SAN
 - configuring, 31
 - unmapping volumes, 113
- SAN volumes, 31, 112, 119, 121, 124, 126
- saving filters, 166
- saving search filters, 166
- sealed mode, 103
- search
 - boolean, 162
 - wildcard, 162
- search filters
 - applying, 166
 - deleting, 166
- search results
 - exporting, 159
 - filtering, 166
 - statistics, 168
- searching log messages, 156
- serial number of SSB, 105
- sharing log files, 128, 133
- shutdown, 84
- Simple Network Management Protocol, 44, 47
- size of a log space, 127
- SMB/CIFS, 63
- SMTP server, 43
- SNMP
 - alerts, 44, 47-48
 - messages, 114
 - queries, 46
 - server, 44
 - SSB MIB, 47
- SNMPv3, 46
- sources, 3, 114
 - creating new, 116
 - defaults, 114
 - SNMP, 114
- spaces, 119
 - creating new, 121
 - indexer delimiters, 123
- split brain, 87-88, 198
- SQL templates, 138
- SSB
 - accounting, 83
 - administrators, 69
 - certificate, 105
 - changelogs, 159
 - configuration (see SSB configuration)
 - configuration changes, 83
 - exporting the configuration of, 97
 - hostname, 40
 - importing the configuration of, 98
 - installation, 204
 - logs, 83
 - reports, 178
 - web certificate, 27
- SSB configuration
 - exporting, 97
 - importing, 98
- SSB options, 149
- SSH
 - console, 100
- SSH connections
 - accessing SSB, 101
- SSH server on SSB, 101
- stable releases, 9
- statistics, 168
 - settings, 177
 - time-based, 177
 - top-least, 177
- status history, 196
- status information via SNMP, 46
- storage, 31
 - installation, 206
- supported browsers, 1, 33
- supported timestamping protocols, 151
- synchronizing data
 - adjusting synchronization speed, 88
- syslog-ng
 - certifications, 153
 - logging configuration changes, 175
 - options, 149
- syslog-ng options, 149
- System monitor, 35
- system monitor
 - number of active hosts, 35
 - number of active senders, 35



system statistics, 196

T

time synchronization, 42

- in HA mode, 42

timestamp, 13, 15

Timestamping Authority

- certificate of, 105

timestamping OID, 151

timestamping protocol, 151

timestamping server, 151

TLS, 7

traceroute, 192

tracking configuration changes, 83

transport layer security (see TLS)

troubleshooting, 192

types of log messages, 156

U

update

- firmware, 93

- in high availability, 95

- license, 96

upgrade

- license, 96

uploading certificates, 108

user groups, 69

user management, 81

- creating usergroups, 79

- finding privileges, 80

- modifying usergroup privileges, 78

- naming usergroups, 80

- searching usergroups, 80

User menu, 34

user preferences, 34

usergroups

- local, 72

users

- web interface, 69, 73

V

volumes, 31, 112, 119, 121, 124, 126

W

web browsers, 33

Welcome Wizard, 23